

Óõ÷íÝò ÅñùôÞóåéò ãéá ôï FreeBSD 6.x êáé 7.x

Ç ìÜäá Ôåêçñßúóçò ôïõ FreeBSD

Íé εΎíáέò Þ òñŬóáέò IBM, AIX, EtherJet, Netfinity, OS/2, PowerPC, PS/2, S/390, έάέ ThinkPad áβίάέ àìðñέέŬ óγíâíεά òçò International Business Machines Corporation óóέò ÇñùíŸíâò Ðíεέòáβâò, Ŭεεâò ÷ þñâò, Þ έάέ óóá äŷí òáòòü÷ñííá.

Íé εΎíáέò IEEE, POSIX, έάέ 802 áβίάέ έáòí÷òñùíŸíá àìðñέέŬ óγíâíεά òñ Institute of Electrical and Electronics Engineers, Inc. óóέò ÇñùíŸíâò Ðíεέòáβâò.

Íé εΎíáέò Intel, Celeron, EtherExpress, i386, i486, Itanium, Pentium, έάέ Xeon áβίάέ àìðñέέŬ óγíâíεά Þ έáòí÷òñùíŸíá àìðñέέŬ óγíâíεά òçò Intel Corporation έάέ òùí òðááòñέέþí òçò óóέò ÇñùíŸíâò Ðíεέòáβâò έάέ óá Ŭεεâò ÷ þñâò.

Íé εΎíáέò Iomega, Zip, έάέ Jaz áβίάέ áβòâ έáòí÷òñùíŸíá àìðñέέŬ óγíâíεά Þ àìðñέέŬ óγíâíεά òçò Iomega Corporation óóέò ÇñùíŸíâò Ðíεέòáβâò Þ/έάέ óá Ŭεεâò ÷ þñâò.

Ŭí Linux áβίάέ Ÿíá έáòí÷òñùíŸíá àìðñέέŬ óγíâíεά òñ Linus Torvalds óóέò ÇñùíŸíâò Ðíεέòáβâò.

Íé εΎíáέò Microsoft, IntelliMouse, MS-DOS, Outlook, Windows, Windows Media, έάέ Windows NT áβίάέ áβòâ έáòí÷òñùíŸíá àìðñέέŬ óγíâíεά Þ àìðñέέŬ óγíâíεά òçò Microsoft Corporation óóέò ÇñùíŸíâò Ðíεέòáβâò έάέ/Þ óá Ŭεεâò ÷ þñâò.

Íé εΎíáέò MIPS έάέ R4000 áβίάέ έáòí÷òñùíŸíá àìðñέέŬ óγíâíεά òçò MIPS Technologies, Inc. óóέò ÇñùíŸíâò Ðíεέòáβâò έάέ óá Ŭεεâò ÷ þñâò.

Íé εΎíáέò Netscape έάέ Netscape Navigator áβίάέ έáòí÷òñùíŸíá àìðñέέŬ óγíâíεά òçò Netscape Communications Corporation óóέò Ç.Ð.Á έάέ Ŭεεâò ÷ þñâò.

Íé εΎíáέò Motif, OSF/1, έάέ UNIX áβίάέ έáòí÷òñùíŸíá àìðñέέŬ óγíâíεά έάέ íé εΎíáέò Þ òñŬóáέò IT DialTone έάέ The Open Group áβίάέ àìðñέέŬ óγíâíεά òñ The Open Group óóέò ÇñùíŸíâò Ðíεέòáβâò έάέ óá Ŭεεâò ÷ þñâò.

Ç εŸíç Oracle áβίάέ έáòí÷òñùíŸíá àìðñέέŬ óγíâíεά òçò Oracle Corporation.

Íé εΎíáέò íé òñŬóáέò Silicon Graphics, SGI, έάέ OpenGL áβίάέ έáòí÷òñùíŸíá àìðñέέŬ óγíâíεά òçò Silicon Graphics, Inc. óóέò ÇñùíŸíâò Ðíεέòáβâò έάέ/Þ óá Ŭεεâò ÷ þñâò óòíí έùóíí.

Íé εΎíáέò Þ òñŬóáέò SPARC, SPARC64, SPARCengine, έάέ UltraSPARC áβίάέ àìðñέέŬ óγíâíεά òçò SPARC International, Inc. óóέò ÇñùíŸíâò Ðíεέòáβâò έάέ óá Ŭεεâò ÷ þñâò. Ç SPARC International Inc. έáòŸ÷áé üεá óá àìðñέέŬ áέέáεþíáòá òùí óòíáüεùí SPARC έάέ áðέòñŸðáé òçí ñèÞ ÷ñÞòç òñò áðü üεá óá íŸέç òçò, έáòüðέí ó÷áðέέÞò Ŭááέáð.

Íé εΎíáέò Þ òñŬóáέò Sun, Sun Microsystems, Java, Java Virtual Machine, JavaServer Pages, JDK, JRE, JSP, JVM, Netra, OpenJDK, Solaris, StarOffice, Sun Blade, Sun Enterprise, Sun Fire, SunOS, Ultra έάέ VirtualBox áβίάέ àìðñέέŬ óγíâíεά Þ έáòí÷òñùíŸíá àìðñέέŬ óγíâíεά òçò Sun Microsystems, Inc. óóέò ÇñùíŸíâò Ðíεέòáβâò έάέ óá Ŭεεâò ÷ þñâò.

Íé εΎíáέò Þ òñŬóáέò U.S. Robotics έάέ Sportster áβίάέ έáòí÷òñùíŸíá àìðñέέŬ óγíâíεά òçò U.S. Robotics Corporation.

Ç εŸíç XFree86 áβίάέ Ÿíá àìðñέέŬ óγíâíεά òñ The XFree86 Project, Inc.

ÐíεεŸð áðü ðέò εΎíáέò Þ òñŬóáέò íé ìðíâò ÷ñçóέííðíέíŸíóáέ áðü òñòò έáðáóέáðáóòŸð Þ òñòò ðùέçòŸð òñòò áέá íá áέáèñβñíóí óá ðñíúúíóá òñòò εáññíŸíóáέ àìðñέέŬ óγíâíεά. ¼ðñò áðóŸð àìðáíβæíóáέ óá áðòü òí έáβíáñ έάέ áέá üóáð áðü áðóŸð áññβæáέ ç ñŬáá ÁíŬðòòíçò òñò FreeBSD üðέ áβίάέ ðέέáñíí íá áβίάέ àìðñέέŬ óγíâíεά, έá ááβòâ Ÿíá áðü óá óγíâíεά: “™” Þ “©”.

ΘΒίαεαò Ðañéα÷ ïñÝíuí

1	ΆέόάũàP	1
2	Ôâèìçñßùóç éáé ÔðíóôPñéιç	7
3	ΆãêáoÜóóáoç	12
4	Óðìááóüöçôá Ôëéêĩý	22
4.1	ÃâíéêÜ	22
4.2	Ãñ÷έóâêôĩíéêÝò éáé ÃðáíãñãáoóÝò	22
4.3	Óêëçñĩß äßóêĩé, ïĩÜãàò ôáéíßàò, ïäçãĩß CD éáé DVD	23
4.4	Ðëçêòññüüãéá éáé ðĩíðßêéá	24
4.5	ÓóóéãòÝò óáéñéáêêðò äðééĩéíüíßàò éáé Äéêôýùóçò	27
4.6	ÓóóéãòÝò P÷ĩõ	28
4.7	¶ëëĩ ðééêü	29
5	Áíôéíáòððéóç ÐñĩäêçìÜòuí	31
6	ÃìðñéêÝò ÃóáññãÝò	46
7	ÃóáññãÝò Ôäééêĩý ×ñPóóç	49
8	Ñýèíéóç ÐòñPíá	53
9	Äßóêĩé, ÓóóòPíáðá Ãñ÷ãßuí éáé ÔñòòòÝò Äêêßççóçò	56
10	Äéá÷ãßñéóç ÓóóòPíáðò	66
11	Ôĩ Óýóóçíá X Windows éáé ïé ÄééñéêÝò Êñóüèàò	74
12	Äéêôýùóç	83
13	ÁóóÜêáéá	89
14	PPP	92
15	ÓáéñéáêÝò Äðééĩéíüíßàò	104
16	ÄéÜöññàò ÁñòòPóáéò	111
17	×éĩýññ éáé FreeBSD	115
18	Ðñĩ÷ũñçìÝíá ÈÝíáðá	118
19	Ãð÷áñéóòßàò	125
	Bibliography	126

ÊáôÛëĩãïò ÐéíÛêùí

3-1. ÌÝãéóôá ïããÝèç áñ÷ãßùí.....	19
12-1. Network cards based on the DEC PCI chipset	85

ΕὰοÛεάέϊ 2 Ôâêìçñßùόç êάέ Òðïóôßñêîç

1. Διέα έάέÛ άέάέßά ððÛñ÷ïïí ó÷άóέέÛ ìά ðï FreeBSD;

Ôï Project δάνÛάέ ìέά ìάñÛέç ãéÛíá ôâêìçñßùόçð ðïð áέάðßéάóάέ online áðü ðïí δάνάέÛðü óýíááóïï: <http://www.FreeBSD.org/docs.html>. Ôá ßάέά áððÛ Ýāñāōā āßíáέ áέάέÝóέíá έάέ ùð δάέÝóά óά ìðïßά ìðñāßōā ìá āāéáóáóóðóāðā āýēēā óðï FreeBSD óýóóçíá óáð. Ðāñéóóüðāñāð éāððñÝñāéāð ó÷άóέέÛ ìά óά δάέÝóά ôâêìçñßùόçð, έá āñāßōā óóéð áéüēïðéāð δāñāñÛïïðð.

Åðéðñüéāðā, έá āñāßōā έάέ Ûēēā óðíéóðñíāíá áέάέßά óðçí Áéāēēñāñāðßā óðï ðÝēïð áððïý ðïð FAQ, έάέ ðïð Åñ÷áéñéāßïð.

2. Åßíáέ áέάέÝóέíç ç ôâêìçñßùόç έάέ óā Ûēēāð ìññðÝð, üðüð áðéü éāßíāñ (ASCII) ð PostScript;

Íáέ. Ç ôâêìçñßùόç āßíáέ áέάέÝóέíç óā ðēðēïð áéāóññāðéēñí ìññðñí έάέ ðñüðñí óðïðßāóçð, óðï áέáéñééóð FTP ðïð FreeBSD, óðïí éáðÛēñāñ /pub/FreeBSD/doc/ (<ftp://ftp.FreeBSD.org/pub/FreeBSD/doc/>).

Ç ôâêìçñßùόç āßíáέ éáðçāññéíðíéçíÝíç ìā áéÛïññïð ðñüðïðð. Áððïß ðāñééāñāÛññïí:

- Ôï üññā ðïð āāññÛïïð, üðüð ð.÷. **faq**, **ð handbook**.
- H āēðóóā έάέ ç èüāéēïðñççç ðïð éāéíÝñð. ÁððÛ āáóßāñíóáέ óóā ìñññāóā ðïð Ý÷ïïí āñēāß óóéð ðïðééÝð ñðēìßóāéð έάέ ðïð ìðñāßōā ìá āñāßōā óðïí éáðÛēñāñ /usr/share/locale óðï FreeBSD óýóóçíÛ óáð. Íé ðñÝ÷ïðóāð āēðóóāð έάέ èüāéēïðñçççð ðïð Ý÷ïïñā áέάέÝóέíāð áððð óç óðéññð óðçí ôâêìçñßùόç āßíáέ ìé δāñāéÛðü:

¼ññā	Åñíçíāßā
en_US.ISO8859-1	ÁāāéééÛ ÇÐÁ
bn_BD.ISO10646-1	Bengali (ð Bangla)
da_DK.ISO8859-1	ÄāíÝæééā
de_DE.ISO8859-1	ÄāññāíééÛ
es_ES.ISO8859-1	ÉóðāíééÛ
fr_FR.ISO8859-1	ÄāēéééÛ
hu_HU.ISO8859-2	ÏðāāñÝæééā
it_IT.ISO8859-15	ÉóāéééÛ
ja_JP.eucJP	ÄéāðñíÝæééā (èüāéēïðñçççç EUC)
mn_MN.UTF-8	ÏāññéééÛ (èüāéēïðñçççç UTF-8)
nl_NL.ISO8859-1	ÏēēāñāéééÛ
pl_PL.ISO8859-2	ÐñēññééÛ
pt_BR.ISO8859-1	ÐñðññāāéééÛ (Äñāæééßā)
ru_RU.KOI8-R	Ñðóééā (èüāéēïðñçççç KOI8-R)
sr_YU.ISO8859-2	ÓÝñāééā
tr_TR.ISO8859-9	Ôñýñēééā
zh_CN.GB2312	ÄðēïðñçççíÝíā ÊéíÝæééā (èüāéēïðñçççç GB2312)
zh_TW.Big5	ÐāñāññóééāÛ ÊéíÝæééā (èüāéēïðñçççç Big5)

6. Δῖο ἰδῖνῃ ἰά ἀνὺ ἀδῖ δῆνῖνῖδ ὀδῖόδῃνείς ἐὰς ἀεὶ βύβος ἀεὶ ὁῖ FreeBSD;

Ὁῖ DaemonNews δῆνῖνῖδ ἰά ἀδῖ δῆνῖνῖδ ἀεὶ βύβος ἐὰς ὀδῖόδῃνείς ἀεὶ ὁῖ FreeBSD. Δῆνῖνῖδ ἰδῖνῖνῖδ ἰδῖνῖνῖδ ἰά ἀνὺ βύβος ὀδῖ ἰδῖνῖνῖδ ὀδῖ BSD Mall (<http://www.bsdmall.com/>).

Ὁῖ FreeBSD Mall δῆνῖνῖδ ἰά ἀδῖ βύβος ἀδῖ δῆνῖνῖδ ὀδῖόδῃνείς ἀεὶ ὁῖ FreeBSD. ἰδῖνῖνῖδ ἰά ἀνὺ βύβος δῆνῖνῖδ ὀδῖ ἰδῖνῖνῖδ ὀδῖ ἰδῖνῖνῖδ ὀδῖ (<http://www.freebsdmall.com/>).

ἰδῖνῖνῖδ ὀδῖ Ὑἱάει ἰδῖνῖνῖδ δῆνῖνῖδ ἰά ἀεὶ βύβος ἐὰς ὀδῖόδῃνείς, ἐὰ δῆνῖνῖδ ἰά ἀεὶ βύβος ἰά ὁῖ Project ἀεὶ ἰά ἐὰς ἰά ἰδῖνῖνῖδ ὀδῖ ὀδῖ ὀδῖ.

Ðáíáëäò 3-1. ÌÝãéóóá ìããÝèç áñ÷ãßùí

ÌÝããèò block fs	ëäéòíõñãß	ðñÝðäé íá ëäéòíõñãß
4K	4T-1	>4T
8K	>32G	32T-1
16K	>128G	32T-1
32K	>512G	64T-1
64K	>2048G	128T-1

¼ðáí òí ÌÝããèò block òíò fs áßíáé 4K, ëäéòíõñãýí ðá ðñéðëÛ Ýíìáóá blocks (triple indirect blocks) éáé ðá ðÛíóá éá Ýðñáðä íá ðñéíñßæííóáé ìùíí áðù òí ÌÝãéóóí áñéèì block ðíò ìðíñãß íá áíáðñáóóáéãß ìä òç ÷ñðóç ðñéðëíí Ýíìáóóíí blocks (ðññßðíò $1K^3 + 1K^2 + 1K$), áëëÛ ðäééëÛ í ðñéíñéóìùð òðãßëáðáé óá Ýíá (ëÛèò) ùñéí 1G-1 óðíòð áñéèìýð òùí blocks. Õí ùñéí óðíòð áñéèìýð òùí block éá Ýðñáðä íá áßíáé 2G-1. ÕðÛñ÷íòí ëÛðíéá ðñíáëðíáðá ùðáí íé áñéèìß òùí block òíò fs ðëçóéÛæíí òí 2G-1, áëëÛ ðÝðíéé áñéèìß block äáí ìðíñýí íá ðñíóããéóóýí ùðáí òí ÌÝããèò block fs áßíáé 4Ê.

Ãéá ìããÝèç block 8Ê éáé ìããéýóãñá, ðá ðÛíóá éá Ýðñáðä íá ðñéíñßæííóáé áðù òí ùñéí 2G-1 óðíòð áñéèìýð òùí block òíò fs, áëëÛ óðçí ðñáñíáðéëùðçðá í ðñéíñéóìùð òðãßëáðáé óðíí ëÛèò ùñéí 1G-1 óðíòð áñéèìýð òùí block òíò fs. H ÷ñðóç òíò óóðóýí ìñßò òùí 2G-1 blocks, äçíéíõñãß ðñÛñíáðé ðñíáëðíáðá.

25. Ãéáðß ðáßñíù òí ìðíóíá ëÛèò, archsw.readin.failed ìäðÛ òçí ìäðáãëðððéóç éáé äëëßçóç íÝíò ððñðíá;

Ãéáðß í ððñðíáð óáð éáé òí ððùéíéðí òíò ááóééý óðóððíáðíò (world) áßíáé äéðùð óðã÷ñíéóíý. Ç ëäéòíõñãßá óá áððð òçí éäðÛóóáóç äáí ððíóðçñßæííóáé. Åããéëùëãðá ùðé ÷ñçóéíðíéãßðá ðéð áíðíëÝð make buildworld éáé make buildkernel äéá íá áíáãéìßðáðá òíí ððñðíá óáð.

Ìðíñãßðá íá äëëéíðóáðá ìñßæííóáð òíí ððñðíá áðððëãßðð áðù òí ääýðãñí óðÛáéí, ðéÝæííóáð ìðíéíððíðá ðëðéðñí ìüéð äãßðá òí ìéáé ðñéí ìäëéíðóáé í loader.

26. Ç äãéáðÛóóáóç éäóãññÝáé éäðÛ òçí äëëßçóç. Õé ìðíñ íá ëÛíù;

ÃíëéíÛóðá íá áðñáññáíðíéðóáðá òçí ððíóððñéíç ACPI. ìüéð ìäëéíðóáé í òííðòððð äëëßçóçð, ðéÝóðá òí ðëðéðñí space. To óýóçíá óáð éá ãíðáíßóáé

OK

. ÅñÛððá

unset acpi_load

éáé éäðùðéí

boot

.

Óçìáépóáéo

1. Óå Ýía e-mail áðü ôĩĩ Keith Frechette <kfrechet@us.ibm.com>.

mixer pcm 100 vol 100 cd 100

4.7 Προετοιμασία

1. Δημιουργία του αρχείου `boot.conf` στο `/boot` του FreeBSD;

Αντίστοιχο του `boot.conf` είναι:

(http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/install.html#INSTALL-MISC) σχετικά με την εγκατάσταση του `boot.conf` στο `/boot`.

2. Ολοκλήρωση του FreeBSD σχετικά με την εγκατάσταση του `boot.conf` στο `/boot`;

Αντίστοιχο του FreeBSD 4.X είναι το `boot.conf`, το οποίο ολοκληρώνεται στο `APM` στο `boot.conf` στο `/boot`. Δημιουργείται το `boot.conf` στο `/boot` με την εντολή `bootconf`.

Αντίστοιχο του FreeBSD 5.X είναι το `boot.conf`, το οποίο ολοκληρώνεται στο `ACPI` στο `boot.conf` στο `/boot`. Δημιουργείται το `boot.conf` στο `/boot` με την εντολή `bootconf`. Αντίστοιχο του `boot.conf` στο `/boot` είναι το `boot.conf` στο `/boot`. Δημιουργείται το `boot.conf` στο `/boot` με την εντολή `bootconf`.

3. Δημιουργία του αρχείου `boot.conf` στο `/boot` του ACPI;

Δημιουργία του αρχείου `boot.conf` στο `/boot` του ACPI:

```
hint.acpi.0.disabled="1"
```

στο αρχείο `/boot/device.hints`.

4. Δημιουργία του Micron στο `boot.conf` στο `/boot`;

Δημιουργία του Micron στο `boot.conf` στο `/boot` του PCI BIOS στο `boot.conf` στο `/boot`. Δημιουργείται το `boot.conf` στο `/boot` με την εντολή `bootconf`.

Αντίστοιχο του `boot.conf` στο `/boot` είναι το `boot.conf` στο `/boot`. Δημιουργείται το `boot.conf` στο `/boot` με την εντολή `bootconf`.

5. Δημιουργία του `boot.conf` στο `/boot` του ASUS K7V. Δημιουργία του `boot.conf` στο `/boot`;

Ολοκλήρωση του BIOS, δημιουργία του `boot.conf` στο `/boot` του `boot virus protection`.

6. Δημιουργία του PCI στο `boot.conf` στο `/boot` του 3Com® στο `boot.conf` στο `/boot` του Micron στο `boot.conf` στο `/boot`;

Δημιουργία του PCI στο `boot.conf` στο `/boot` του Micron στο `boot.conf` στο `/boot` του PCI BIOS στο `boot.conf` στο `/boot`. Δημιουργείται το `boot.conf` στο `/boot` με την εντολή `bootconf`.

Αντίστοιχο του `boot.conf` στο `/boot` είναι το `boot.conf` στο `/boot`. Δημιουργείται το `boot.conf` στο `/boot` με την εντολή `bootconf`.

ΕὰöÜëáéĩ 6 ÅìðĩñééÝò ÅöáñĩĩãÝò

Όçĩåßüόç: This section is still very sparse, though we are hoping, of course, that companies will add to it! :) The FreeBSD group has no financial interest in any of the companies listed here but simply lists them as a public service (and feels that commercial interest in FreeBSD can have very positive effects on FreeBSD's long-term viability). We encourage commercial software vendors to send their entries here for inclusion. See the Vendors page (<http://www.FreeBSD.org/commercial/index.html>) for a longer list.

1. Where can I get an Office Suite for FreeBSD?

The open-source OpenOffice.org (<http://www.openoffice.org>) office suite works natively on FreeBSD. The Linux version of StarOffice (<http://www.sun.com/staroffice/>), the value-added closed-source version of OpenOffice.org, also works on FreeBSD.

FreeBSD also includes a variety of text editors, spreadsheets, and drawing programs in the Ports Collection.

2. Where can I get Motif® for FreeBSD?

The Open Group has released the source code to Motif 2.2.2. You can install the `open-motif` package, or compile it from ports. Refer to the ports section of the Handbook

(http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/ports.html) for more information on how to do this.

Όçĩåßüόç: The Open Motif distribution only allows redistribution if it is running on an open source (<http://www.opensource.org/>) operating system.

In addition, there are commercial distributions of the Motif software available. These, however, are not for free, but their license allows them to be used in closed-source software. Contact Apps2go for the least expensive ELF Motif 2.1.20 distribution for FreeBSD (either i386 or Alpha).

There are two distributions, the “development edition” and the “runtime edition” (for much less). These distributions includes:

- OSF/Motif manager, xmbind, panner, wsm.
- Development kit with uil, mrm, xm, xmcxx, include and Imake files.
- Static and dynamic ELF libraries.
- Demonstration applets.

Be sure to specify that you want the FreeBSD version of Motif when ordering (do not forget to mention the architecture you want too)! Versions for NetBSD and OpenBSD are also sold by *Apps2go*. This is currently a FTP only download.

More info

Apps2go WWW page (<http://www.apps2go.com/>)

or

<sales@apps2go.com> or <support@apps2go.com>

or

phone (817) 431 8775 or +1 817 431-8775

Contact Xi Graphics for an a.out Motif 2.0 distribution for FreeBSD.

This distribution includes:

- OSF/Motif manager, xmbind, panner, wsm.
- Development kit with uil, mrm, xm, xmcxx, include and Imake files.
- Static and dynamic libraries (for use with FreeBSD 2.2.8 and earlier).
- Demonstration applets.
- Preformatted manual pages.

Be sure to specify that you want the FreeBSD version of Motif when ordering! Versions for BSDI and Linux are also sold by *Xi Graphics*. This is currently a 4 diskette set... in the future this will change to a unified CD distribution like their CDE.

3. Where can I get CDE for FreeBSD?

Xi Graphics used to sell CDE for FreeBSD, but no longer do.

KDE (<http://www.kde.org/>) is an open source X11 desktop which is similar to CDE in many respects. You might also like the look and feel of xfce (<http://www.xfce.org/>). KDE and xfce are both in the ports system (<http://www.FreeBSD.org/ports/index.html>).

4. Are there any commercial high-performance X servers?

Yes, Xi Graphics (<http://www.xig.com/>) sells Accelerated-X products for FreeBSD and other Intel based systems.

The Xi Graphics offering is a high performance X Server that offers easy configuration, support for multiple concurrent video boards and is distributed in binary form only, in a unified diskette distribution for FreeBSD and Linux. Xi Graphics also offers a high performance X Server tailored for laptop support.

There is a free “compatibility demo” of version 5.0 available.

Xi Graphics also sells Motif and CDE for FreeBSD (see above).

More info

Xi Graphics WWW page (<http://www.xig.com/>)

or

<sales@xig.com> or <support@xig.com>

or

phone (800) 946 7433 or +1 303 298-7478.

5. Are there any Database systems for FreeBSD?

Yes! See the Commercial Vendors

(http://www.FreeBSD.org/commercial/software_bycat.html#CATEGORY_DATABASE) section of FreeBSD's Web site.

Also see the Databases (<http://www.FreeBSD.org/ports/databases.html>) section of the Ports collection.

6. Can I run Oracle® on FreeBSD?

Yes. The following pages tell you exactly how to set up Linux-Oracle on FreeBSD:

- <http://www.unixcities.com/oracle/index.html> (<http://www.unixcities.com/oracle/index.html>)
- <http://www.shadowcom.net/freebsd-oracle9i/> (<http://www.shadowcom.net/freebsd-oracle9i/>)

ΕὰöÜëáéï 7 ÅöáñïïãÝò Ôãëëëïÿ × ñÞóôç

1. So, where are all the user applications?

Please take a look at the ports page (<http://www.FreeBSD.org/ports/index.html>) for info on software packages ported to FreeBSD. The list currently tops 24,000 and is growing daily, so come back to check often or subscribe to the `freebsd-announce` mailing list for periodic updates on new entries.

Most ports should work on the 4.X, 5.X, and 6.X branches. Each time a FreeBSD release is made, a snapshot of the ports tree at the time of release is also included in the `ports/` directory.

We also support the concept of a “package”, essentially no more than a compressed binary distribution with a little extra intelligence embedded in it for doing whatever custom installation work is required. A package can be installed and uninstalled again easily without having to know the gory details of which files it includes.

Use the package installation menu in `/stand/sysinstall` (under the post-configuration menu item) or invoke the `pkg_add(1)` command on the specific package files you are interested in installing. Package files can usually be identified by their `.tgz` or `.tbz` suffix and CDROM distribution people will have a `packages/All` directory on their CD which contains such files. They can also be downloaded over the net for various versions of FreeBSD at the following locations:

for 4.X-RELEASE/4-STABLE

`ftp://ftp.FreeBSD.org/pub/FreeBSD/ports/i386/packages-4-stable/`
(`ftp://ftp.FreeBSD.org/pub/FreeBSD/ports/i386/packages-4-stable/`)

for 5.X-RELEASE/5-STABLE

`ftp://ftp.FreeBSD.org/pub/FreeBSD/ports/i386/packages-5-stable`
(`ftp://ftp.FreeBSD.org/pub/FreeBSD/ports/i386/packages-5-stable/`)

for 6.X-RELEASE/6-STABLE

`ftp://ftp.FreeBSD.org/pub/FreeBSD/ports/i386/packages-6-stable`
(`ftp://ftp.FreeBSD.org/pub/FreeBSD/ports/i386/packages-6-stable/`)

for 7-CURRENT

`ftp://ftp.FreeBSD.org/pub/FreeBSD/ports/i386/packages-7-current`
(`ftp://ftp.FreeBSD.org/pub/FreeBSD/ports/i386/packages-7-current/`)

or your nearest local mirror site.

Note that all ports may not be available as packages since new ones are constantly being added. It is always a good idea to check back periodically to see which packages are available at the `ftp.FreeBSD.org` (`ftp://ftp.FreeBSD.org/pub/FreeBSD/`) master site.

2. How do I configure INN (Internet News) for my machine?

After installing the `news/inn` package or port, an excellent place to start is Dave Barr’s INN Page (<http://www.visi.com/~barr/INN.html>) where you will find the INN FAQ.

3. Does FreeBSD support Java™?

Yes. Please see <http://www.FreeBSD.org/java/> (<http://www.FreeBSD.org/java/index.html>).

4. Why can I not build this port on my 4.X-STABLE machine?

If you are running a FreeBSD version that lags significantly behind -CURRENT or -STABLE, you may need to update your ports collection; see the [Keeping Up](http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/porters-handbook/keeping-up.html) (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/porters-handbook/keeping-up.html) section of the Porter's Handbook for further information on how to do this. If you are up to date, then someone might have committed a change to the port which works for -CURRENT but which broke the port for -STABLE. Please submit a bug report on this with the `send-pr(1)` command, since the ports collection is supposed to work for both the -CURRENT and -STABLE branches.

5. I just tried to build INDEX using `make index`, and it failed. Why?

First, always make sure that you have a completely up-to-date Ports Collection. Errors that affect building INDEX from an up-to-date copy of the Ports Collection are high-visibility and are thus almost always fixed immediately.

However, if you are up-to-date, perhaps you are seeing another problem. `make index` has a known bug in dealing with incomplete copies of the Ports Collection. It assumes that you have a local copy of every single port that every other port that you have a local copy of depends on. To explain, if you have a copy of `foo/bar` on your disk, and `foo/bar` depends on `baz/quux`, then you must also have a copy of `baz/quux` on your disk, and the ports `baz/quux` depends on, and so on. Otherwise, `make index` has insufficient information to create its dependency tree.

This is particularly a problem for FreeBSD users who utilize `cvsup(1)` to track the Ports Collection but choose not to install certain categories by specifying them in `refuse`. In theory, one should be able to refuse categories, but in practice there are too many ports that depend on ports in other categories. Until someone comes up with a solution for this problem, the general rule is that if you want to build INDEX, you must have a complete copy of the Ports Collection.

There are rare cases where INDEX will not build due to odd cases involving `WITH_*` or `WITHOUT_*` variables being set in `make.conf`. If you suspect that this is the case, please try to make INDEX with those Makevars turned off before reporting it to freebsd-ports@lists.FreeBSD.org (<http://lists.FreeBSD.org/mailman/listinfo/freebsd-ports>).

6. Why is CVSup not integrated in the main FreeBSD tree?

The FreeBSD base system is designed as self-hosting - it should be possible to build the whole operating system starting with a very limited set of tools. Thus, the actual build tools needed to compile the FreeBSD sources are bundled with the sources themselves. This includes a C compiler (`gcc(1)`), `make(1)`, `awk(1)`, and similar tools.

Since CVSup is written in Modula-3, adding it to the FreeBSD base system would also require adding and maintaining a Modula-3 compiler. This would lead to both an increase in the disk space consumed by the FreeBSD sources and additional maintenance work. Thus, it is much easier for both the developers and users to keep CVSup as a separate port, which can be easily installed as a package bundled on the FreeBSD installation CDs.

7. I updated the sources, now how do I update my installed ports?

FreeBSD does not include a port upgrading tool, but it does have some tools to make the upgrade process somewhat easier. You can also install additional tools to simplify port handling.

The `pkg_version(1)` command can generate a script that will update installed ports to the latest version in the ports tree.

```
# pkg_version -c > /tmp/myscript
```

The output script *must* be edited by hand before you use it. Recent versions of `pkg_version(1)` force this by inserting an `exit(1)` at the beginning of the script.

You should save the output of the script, as it will note packages that depend on the one that has been updated. These may or may not need to be updated as well. The usual case where they need to be updated is that a shared library has changed version numbers, so the ports that used that library need to be rebuilt to use the new version.

Όχι ἀβουόç: Beginning with FreeBSD 5.0 (and higher revisions), `pkg_version(1)` no longer supports the `-c` option.

If you have the disk space, you can use the `portupgrade` tool to automate all of this. `portupgrade` includes various tools to simplify package handling. It is available under `ports-mgmt/portupgrade`. Since it is written in Ruby, `portupgrade` is an unlikely candidate for integration with the main FreeBSD tree. That should not stop anyone from using it, however.

If your system is up full time, the `periodic(8)` system can be used to generate a weekly list of ports that might need updating by setting `weekly_status_pkg_enable="YES"` in `/etc/periodic.conf`.

8. Why is `/bin/sh` so minimal? Why does FreeBSD not use `bash` or another shell?

Because POSIX® says that there shall be such a shell.

The more complicated answer: many people need to write shell scripts which will be portable across many systems. That is why POSIX specifies the shell and utility commands in great detail. Most scripts are written in Bourne shell, and because several important programming interfaces (`make(1)`, `system(3)`, `popen(3)`, and analogues in higher-level scripting languages like Perl and Tcl) are specified to use the Bourne shell to interpret commands. Because the Bourne shell is so often and widely used, it is important for it to be quick to start, be deterministic in its behavior, and have a small memory footprint.

The existing implementation is our best effort at meeting as many of these requirements simultaneously as we can. In order to keep `/bin/sh` small, we have not provided many of the convenience features that other shells have. That is why the Ports Collection includes more featureful shells like `bash`, `scsh`, `tcsh`, and `zsh`. (You can compare for yourself the memory utilization of all these shells by looking at the “VSZ” and “RSS” columns in a `ps -u` listing.)

9. Why do Netscape and Opera take so long to start?

The usual answer is that DNS on your system is misconfigured. Both Netscape and Opera perform DNS checks when starting up. The browser will not appear on your desktop until the program either gets a response or determines that the system has no network connection.

10. I updated parts of the Ports Collection using CVSup, and now many ports fail to build with mysterious error messages! What happened? Is the Ports Collection broken in some major way?

If you only update parts of the Ports Collection, using one of its CVSup subcollections and not the `ports-all` CVSup collection, you should *always* update the `ports-base` subcollection too! The reasons are described in the

Handbook

(http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/cvsup.html#CVSUP-COLLEC-PBASE-WARN).

11. How do I create audio CDs from my MIDI files?

To create audio CDs from MIDI files, first install `audio/timidity++` from ports then install manually the GUS patches set by Eric A. Welsh, available at <http://www.stardate.bc.ca/eawpatches/html/default.htm>. After `timidity++` has been installed properly, midi files may be converted to wav files with the following command line:

```
% timidity -Ow -s 44100 -o /tmp/juke/01.wav 01.mid
```

The wav files can then be converted to other formats or burned onto audio CDs, as described in the FreeBSD Handbook.

ΕὰöÜëáéï 8 Ñýèìéόç ÐõñΠία

1. I would like to customize my kernel. Is it difficult?

Not at all! Check out the `kernel config` section of the Handbook

(http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/kernelconfig.html).

ΌçìäΒùόç: We recommend that you make a dated snapshot of your new `/kernel` called `/kernel.YYMMDD` after you get it working properly. Also back up your new `/modules` directory to `/modules.YYMMDD`. That way, if you make a mistake the next time you play with your configuration you can boot the backup kernel instead of having to fall back to `kernel.GENERIC`. This is particularly important if you are now booting from a controller that `GENERIC` does not support.

2. My kernel compiles fail because `_hw_float` is missing. How do I solve this problem?

You probably removed `npx0` (see `npx(4)`) from your kernel configuration file because you do not have a math co-processor. The `npx0` device is *MANDATORY*. Somewhere inside your hardware lies a device that provides hardware floating-point support, even if it is no longer a separate device as used in the good old 386 days. You *must* include the `npx0` device. Even if you manage to build a kernel without `npx0` support, it will not boot anyway.

3. Why is my kernel so big (over 10MB)?

Chances are, you compiled your kernel in *debug mode*. Kernels built in debug mode contain many symbols that are used for debugging, thus greatly increasing the size of the kernel. Note that there will be little or no performance decrease from running a debug kernel, and it is useful to keep one around in case of a system panic.

However, if you are running low on disk space, or you simply do not want to run a debug kernel, make sure that both of the following are true:

- You do not have a line in your kernel configuration file that reads:
`makeoptions DEBUG=-g`
- You are not running `config(8)` with the `-g` option.

Either of the above settings will cause your kernel to be built in debug mode. As long as you make sure you follow the steps above, you can build your kernel normally, and you should notice a fairly large size decrease; most kernels tend to be around 1.5MB to 2MB.

4. Why do I get interrupt conflicts with multi-port serial code?

When I compile a kernel with multi-port serial code, it tells me that only the first port is probed and the rest skipped due to interrupt conflicts. How do I fix this?

The problem here is that FreeBSD has code built-in to keep the kernel from getting trashed due to hardware or software conflicts. The way to fix this is to leave out the IRQ settings on all but one port. Here is an example:

```
#
```

```
# Multiport high-speed serial line - 16550 UARTS
#
device sio2 at isa? port 0x2a0 tty irq 5 flags 0x501 vector siointr
device sio3 at isa? port 0x2a8 tty flags 0x501 vector siointr
device sio4 at isa? port 0x2b0 tty flags 0x501 vector siointr
device sio5 at isa? port 0x2b8 tty flags 0x501 vector siointr
```

5. Why does every kernel I try to build fail to compile, even GENERIC?

There are a number of possible causes for this problem. They are, in no particular order:

- You are not using the new `make buildkernel` and `make installkernel` targets, and your source tree is different from the one used to build the currently running system (e.g., you are compiling 4.3-RELEASE on a 4.0-RELEASE system). If you are attempting an upgrade, please read the `/usr/src/UPDATING` file, paying particular attention to the “COMMON ITEMS” section at the end.
- You are using the new `make buildkernel` and `make installkernel` targets, but you failed to assert the completion of the `make buildworld` target. The `make buildkernel` target relies on files generated by the `make buildworld` target to complete its job correctly.
- Even if you are trying to build FreeBSD-STABLE, it is possible that you fetched the source tree at a time when it was either being modified, or broken for other reasons; only releases are absolutely guaranteed to be buildable, although FreeBSD-STABLE builds fine the majority of the time. If you have not already done so, try re-fetching the source tree and see if the problem goes away. Try using a different server in case the one you are using is having problems.

6. How can I verify which scheduler is in use on a running system?

If you are running FreeBSD version 5.2.1 or earlier, check for the existence of the `kern.quantum` sysctl. If you have it, you should see something like this:

```
% sysctl kern.quantum
kern.sched.quantum: 99960
```

If the `kern.quantum` sysctl exists, you are using the 4BSD scheduler. If not, you will get an error printed by `sysctl(8)` (which you can safely ignore):

```
% sysctl kern.sched.quantum
sysctl: unknown oid 'kern.sched.quantum'
```

In FreeBSD version 5.3-RELEASE and later, the name of the scheduler currently being used is directly available as the value of the `kern.sched.name` sysctl:

```
% sysctl kern.sched.name
kern.sched.name: 4BSD
```

7. What is `kern.quantum`?

`kern.quantum` is the maximum number of ticks a process can run without being preempted. It is specific to the 4BSD scheduler, so you can use its presence or absence to determine which scheduler is in use. In FreeBSD 5.X or later `kern.quantum` has been renamed to `kern.sched.quantum`.

8. What is `kern.sched.quantum`?

See ¶ 7.

ΕὰöÛëáéĩ 9 ÄΒóêïé, ÓõóôΠιάôá Äñ÷åΒùí êáé ÖïñôùôÝò ÄêêΒίçóçò

1. How do I move my system over to my huge new disk?

The best way is to reinstall the OS on the new disk, then move the user data over. This is highly recommended if you have been tracking -STABLE for more than one release, or have updated a release instead of installing a new one. You can install booteasy on both disks with `boot0cfg(8)`, and dual boot them until you are happy with the new configuration. Skip the next paragraph to find out how to move the data after doing this.

Should you decide not to do a fresh install, you need to partition and label the new disk with either `/stand/sysinstall`, or `fdisk(8)` and `disklabel(8)`. You should also install booteasy on both disks with `boot0cfg(8)`, so that you can dual boot to the old or new system after the copying is done.

Now you have the new disk set up, and are ready to move the data. Unfortunately, you cannot just blindly copy the data. Things like device files (in `/dev`), flags, and links tend to screw that up. You need to use tools that understand these things, which means `dump(8)`. Although it is suggested that you move the data in single user mode, it is not required.

You should never use anything but `dump(8)` and `restore(8)` to move the root filesystem. The `tar(1)` command may work - then again, it may not. You should also use `dump(8)` and `restore(8)` if you are moving a single partition to another empty partition. The sequence of steps to use `dump` to move a partitions data to a new partition is:

1. `newfs` the new partition.
2. mount it on a temporary mount point.
3. `cd` to that directory.
4. `dump` the old partition, piping output to the new one.

For example, if you are going to move root to `/dev/ad1s1a`, with `/mnt` as the temporary mount point, it is:

```
# newfs /dev/ad1s1a
# mount /dev/ad1s1a /mnt
# cd /mnt
# dump 0af - / | restore xf -
```

Rearranging your partitions with `dump` takes a bit more work. To merge a partition like `/var` into its parent, create the new partition large enough for both, move the parent partition as described above, then move the child partition into the empty directory that the first move created:

```
# newfs /dev/ad1s1a
# mount /dev/ad1s1a /mnt
# cd /mnt
# dump 0af - / | restore xf -
# cd var
# dump 0af - /var | restore xf -
```

To split a directory from its parent, say putting `/var` on its own partition when it was not before, create both partitions, then mount the child partition on the appropriate directory in the temporary mount point, then move the old single partition:

```
# newfs /dev/ad1s1a
# newfs /dev/ad1s1d
# mount /dev/ad1s1a /mnt
# mkdir /mnt/var
# mount /dev/ad1s1d /mnt/var
# cd /mnt
# dump 0af - / | restore xf -
```

You might prefer `cpio(1)`, `pax(1)`, `tar(1)` to `dump(8)` for user data. At the time of this writing, these are known to lose file flag information, so use them with caution.

2. Will a “dangerously dedicated” disk endanger my health?

The installation procedure allows you to chose two different methods in partitioning your hard disk(s). The default way makes it compatible with other operating systems on the same machine, by using fdisk table entries (called “slices” in FreeBSD), with a FreeBSD slice that employs partitions of its own. Optionally, one can chose to install a boot-selector to switch between the possible operating systems on the disk(s). The alternative uses the entire disk for FreeBSD, and makes no attempt to be compatible with other operating systems.

So why it is called “dangerous”? A disk in this mode does not contain what normal PC utilities would consider a valid fdisk table. Depending on how well they have been designed, they might complain at you once they are getting in contact with such a disk, or even worse, they might damage the BSD bootstrap without even asking or notifying you. In addition, the “dangerously dedicated” disk’s layout is known to confuse many BIOSes, including those from AWARD (e.g. as found in HP Netserver and Micronics systems as well as many others) and Symbios/NCR (for the popular 53C8xx range of SCSI controllers). This is not a complete list, there are more. Symptoms of this confusion include the `read error` message printed by the FreeBSD bootstrap when it cannot find itself, as well as system lockups when booting.

Why have this mode at all then? It only saves a few kbytes of disk space, and it can cause real problems for a new installation. “Dangerously dedicated” mode’s origins lie in a desire to avoid one of the most common problems plaguing new FreeBSD installers - matching the BIOS “geometry” numbers for a disk to the disk itself.

“Geometry” is an outdated concept, but one still at the heart of the PC’s BIOS and its interaction with disks. When the FreeBSD installer creates slices, it has to record the location of these slices on the disk in a fashion that corresponds with the way the BIOS expects to find them. If it gets it wrong, you will not be able to boot.

“Dangerously dedicated” mode tries to work around this by making the problem simpler. In some cases, it gets it right. But it is meant to be used as a last-ditch alternative - there are better ways to solve the problem 99 times out of 100.

So, how do you avoid the need for “DD” mode when you are installing? Start by making a note of the geometry that your BIOS claims to be using for your disks. You can arrange to have the kernel print this as it boots by specifying `-v` at the `boot:` prompt, or using `boot -v` in the loader. Just before the installer starts, the kernel will print a list of BIOS geometries. Do not panic - wait for the installer to start and then use scrollbar to read the numbers. Typically the BIOS disk units will be in the same order that FreeBSD lists your disks, first IDE, then SCSI.

When you are slicing up your disk, check that the disk geometry displayed in the FDISK screen is correct (ie. it matches the BIOS numbers); if it is wrong, use the `g` key to fix it. You may have to do this if there is absolutely

nothing on the disk, or if the disk has been moved from another system. Note that this is only an issue with the disk that you are going to boot from; FreeBSD will sort itself out just fine with any other disks you may have.

Once you have got the BIOS and FreeBSD agreeing about the geometry of the disk, your problems are almost guaranteed to be over, and with no need for “DD” mode at all. If, however, you are still greeted with the dreaded `read error` message when you try to boot, it is time to cross your fingers and go for it - there is nothing left to lose.

To return a “dangerously dedicated” disk for normal PC use, there are basically two options. The first is, you write enough NULL bytes over the MBR to make any subsequent installation believe this to be a blank disk. You can do this for example with

```
# dd if=/dev/zero of=/dev/rda0 count=15
```

Alternatively, the undocumented DOS “feature”

```
C:\> fdisk /mbr
```

will install a new master boot record as well, thus clobbering the BSD bootstrap.

3. Which partitions can safely use Soft Updates? I have heard that Soft Updates on / can cause problems.

Short answer: you can usually use Soft Updates safely on all partitions.

Long answer: There used to be some concern over using Soft Updates on the root partition. Soft Updates has two characteristics that caused this. First, a Soft Updates partition has a small chance of losing data during a system crash. (The partition will not be corrupted; the data will simply be lost.) Also, Soft Updates can cause temporary space shortages.

When using Soft Updates, the kernel can take up to thirty seconds to actually write changes to the physical disk. If you delete a large file, the file still resides on disk until the kernel actually performs the deletion. This can cause a very simple race condition. Suppose you delete one large file and immediately create another large file. The first large file is not yet actually removed from the physical disk, so the disk might not have enough room for the second large file. You get an error that the partition does not have enough space, although you know perfectly well that you just released a large chunk of space! When you try again mere seconds later, the file creation works as you expect. This has left more than one user scratching his head and doubting his sanity, the FreeBSD filesystem, or both.

If a system should crash after the kernel accepts a chunk of data for writing to disk, but before that data is actually written out, data could be lost or corrupted. This risk is extremely small, but generally manageable. Use of IDE write caching greatly increases this risk; it is strongly recommended that you disable IDE write caching when using Soft Updates.

These issues affect all partitions using Soft Updates. So, what does this mean for the root partition?

Vital information on the root partition changes very rarely. Files such as `/kernel` and the contents of `/etc` only change during system maintenance, or when users change their passwords. If the system crashed during the thirty-second window after such a change is made, it is possible that data could be lost. This risk is negligible for most applications, but you should be aware that it exists. If your system cannot tolerate this much risk, do not use Soft Updates on the root filesystem!

`/` is traditionally one of the smallest partitions. By default, FreeBSD puts the `/tmp` directory on `/`. If you have a busy `/tmp`, you might see intermittent space problems. Symlinking `/tmp` to `/var/tmp` will solve this problem.

4. What is inappropriate about my ccd?

The symptom of this is:

```
# ccdconfig -C
ccdconfig: ioctl (CCDIOCSET): /dev/ccd0c: Inappropriate file type or format
```

This usually happens when you are trying to concatenate the `c` partitions, which default to type `unused`. The `ccd` driver requires the underlying partition type to be `FS_BSDFFS`. Edit the disklabel of the disks you are trying to concatenate and change the types of partitions to `4.2BSD`.

5. Why can I not edit the disklabel on my ccd?

The symptom of this is:

```
# disklabel ccd0
(it prints something sensible here, so let us try to edit it)
# disklabel -e ccd0
(edit, save, quit)
disklabel: ioctl DIOWDINFO: No disk label on disk;
use "disklabel -r" to install initial label
```

This is because the disklabel returned by `ccd` is actually a “fake” one that is not really on the disk. You can solve this problem by writing it back explicitly, as in:

```
# disklabel ccd0 > /tmp/disklabel.tmp
# disklabel -Rr ccd0 /tmp/disklabel.tmp
# disklabel -e ccd0
(this will work now)
```

6. Can I mount other foreign filesystems under FreeBSD?

FreeBSD supports a variety of other filesystems.

Digital UNIX

UFS CDRoms can be mounted directly on FreeBSD. Mounting disk partitions from Digital UNIX and other systems that support UFS may be more complex, depending on the details of the disk partitioning for the operating system in question.

Linux

FreeBSD supports `ext2fs` partitions. See `mount_ext2fs(8)` for more information.

Windows NT®

FreeBSD includes a read-only NTFS driver. For more information, see `mount_ntfs(8)`.

FAT

FreeBSD includes a read-write FAT driver. For more information, see `mount_msdosfs(8)`.

ReiserFS

FreeBSD includes a read-only ReiserFS driver. For more information, see `mount_reiserfs(8)`.

FreeBSD also supports network filesystems such as NFS (see `mount_nfs(8)`), NetWare (see `mount_nwfs(8)`), and Microsoft-style SMB filesystems (see `mount_smbfs(8)`).

7. How do I mount a secondary DOS partition?

The secondary DOS partitions are found after ALL the primary partitions. For example, if you have an “E” partition as the second DOS partition on the second SCSI drive, you need to create the special files for “slice 5” in `/dev`, then `mount /dev/dals5:`

```
# cd /dev
# sh MAKEDEV dals5
# mount -t msdosfs /dev/dals5 /dos/e
```

ὈçìἈΒὺòç: You can omit this step if you are running FreeBSD 5.0-RELEASE or newer with `devfs(5)` enabled.

8. Is there a cryptographic filesystem for FreeBSD?

Yes. FreeBSD 5.0 includes `gbde(8)`, and FreeBSD 6.0 added `geli(8)`. For earlier releases, see the `security/cfs` port.

9. How can I use the Windows NT loader to boot FreeBSD?

The general idea is that you copy the first sector of your native root FreeBSD partition into a file in the DOS/Windows NT partition. Assuming you name that file something like `c:\bootsect.bsd` (inspired by `c:\bootsect.dos`), you can then edit the `c:\boot.ini` file to come up with something like this:

```
[boot loader]
timeout=30
default=multi(0)disk(0)rdisk(0)partition(1)\WINDOWS
[operating systems]
multi(0)disk(0)rdisk(0)partition(1)\WINDOWS="Windows NT"
C:\BOOTSECT.BSD="FreeBSD"
C:\="DOS"
```

If FreeBSD is installed on the same disk as the Windows NT boot partition simply copy `/boot/boot1` to `C:\BOOTSECT.BSD`. However, if FreeBSD is installed on a different disk `/boot/boot1` will not work, `/boot/boot0` is needed.

`/boot/boot0` needs to be installed using `sysinstall` by selecting the FreeBSD boot manager on the screen which asks if you wish to use a boot manager. This is because `/boot/boot0` has the partition table area filled with NULL characters but `sysinstall` copies the partition table before copying `/boot/boot0` to the MBR.

ḐñìᾶᾷòἸḐçόç: Do not simply copy `/boot/boot0` instead of `/boot/boot1`; you will overwrite your partition table and render your computer un-bootable!

When the FreeBSD boot manager runs it records the last OS booted by setting the active flag on the partition table entry for that OS and then writes the whole 512-bytes of itself back to the MBR so if you just copy `/boot/boot0` to `C:\BOOTSECT.BSD` then it writes an empty partition table, with the active flag set on one entry, to the MBR.

10. How do I boot FreeBSD and Linux from LILO?

If you have FreeBSD and Linux on the same disk, just follow LILO's installation instructions for booting a non-Linux operating system. Very briefly, these are:

Boot Linux, and add the following lines to `/etc/lilo.conf`:

```
other=/dev/hda2
    table=/dev/hda
    label=FreeBSD
```

(the above assumes that your FreeBSD slice is known to Linux as `/dev/hda2`; tailor to suit your setup). Then, run `lilo` as `root` and you should be done.

If FreeBSD resides on another disk, you need to add `loader=/boot/chain.b` to the LILO entry. For example:

```
other=/dev/dab4
    table=/dev/dab
    loader=/boot/chain.b
    label=FreeBSD
```

In some cases you may need to specify the BIOS drive number to the FreeBSD boot loader to successfully boot off the second disk. For example, if your FreeBSD SCSI disk is probed by BIOS as BIOS disk 1, at the FreeBSD boot loader prompt you need to specify:

```
Boot: 1:da(0,a)/kernel
```

You can configure `boot(8)` to automatically do this for you at boot time.

The `Linux+FreeBSD mini-HOWTO` (<http://sunsite.unc.edu/LDP/HOWTO/mini/Linux+FreeBSD.html>) is a good reference for FreeBSD and Linux interoperability issues.

11. How do I boot FreeBSD and Linux using GRUB

Booting FreeBSD using GRUB is very simple. Just add the following to your configuration file `/boot/grub/grub.conf`.

```
title FreeBSD 6.1
    root (hd0,a)
    kernel /boot/loader
```

Where `hd0, a` points to your root partition on the first disk. If you need to specify which slice number should be used, use something like this `(hd0, 2, a)`. By default, if the slice number is omitted, GRUB searches the first slice which has 'a' partition.

12. How do I boot FreeBSD and Linux using BootEasy?

Install LILO at the start of your Linux boot partition instead of in the Master Boot Record. You can then boot LILO from BootEasy.

If you are running Windows 95 and Linux this is recommended anyway, to make it simpler to get Linux booting again if you should need to reinstall Windows 95 (which is a Jealous Operating System, and will bear no other Operating Systems in the Master Boot Record).

13. How do I change the boot prompt from ??? to something more meaningful?

You can not do that with the standard boot manager without rewriting it. There are a number of other boot managers in the `sysutils` ports category that provide this functionality.

14. I have a new removable drive, how do I use it?

Whether it is a removable drive like a Zip or an EZ drive (or even a floppy, if you want to use it that way), or a new hard disk, once it is installed and recognized by the system, and you have your cartridge/floppy/whatever slotted in, things are pretty much the same for all devices.

(this section is based on Mark Mayo's ZIP FAQ (<http://www.vmunix.com/mark/FreeBSD/ZIP-FAQ.html>))

If it is a ZIP drive or a floppy, you have already got a DOS filesystem on it, you can use a command like this:

```
# mount -t msdosfs /dev/fd0c /floppy
```

if it is a floppy, or this:

```
# mount -t msdosfs /dev/da2s4 /zip
```

for a ZIP disk with the factory configuration.

For other disks, see how they are laid out using `fdisk(8)` or `sysinstall(8)`.

The rest of the examples will be for a ZIP drive on da2, the third SCSI disk.

Unless it is a floppy, or a removable you plan on sharing with other people, it is probably a better idea to stick a BSD filesystem on it. You will get long filename support, at least a 2X improvement in performance, and a lot more stability. First, you need to redo the DOS-level partitions/filesystems. You can either use `fdisk(8)` or `/stand/sysinstall`, or for a small drive that you do not want to bother with multiple operating system support on, just blow away the whole FAT partition table (slices) and just use the BSD partitioning:

```
# dd if=/dev/zero of=/dev/rda2 count=2
# disklabel -Brw da2 auto
```

You can use `disklabel` or `/stand/sysinstall` to create multiple BSD partitions. You will certainly want to do this if you are adding swap space on a fixed disk, but it is probably irrelevant on a removable drive like a ZIP.

Finally, create a new filesystem, this one is on our ZIP drive using the whole disk:

```
# newfs /dev/rda2c
```

and mount it:

```
# mount /dev/da2c /zip
```

and it is probably a good idea to add a line like this to `/etc/fstab` (see `fstab(5)`) so you can just type `mount /zip` in the future:

```
/dev/da2c /zip ffs rw,noauto 0 0
```

15. Why do I get `Incorrect super block` when mounting a CDROM?

You have to tell `mount(8)` the type of the device that you want to mount. This is described in the Handbook section on optical media (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/creating-cds.html), specifically the section Using Data CDs (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/creating-cds.html#MOUNTING-CD).

16. Why do I get `Device not configured` when mounting a CDROM?

This generally means that there is no CDROM in the CDROM drive, or the drive is not visible on the bus. Please see the Using Data CDs (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/creating-cds.html#MOUNTING-CD) section of the Handbook for a detailed discussion of this issue.

17. Why do all non-English characters in filenames show up as “?” on my CDs when mounted in FreeBSD?

Your CDROM probably uses the “Joliet” extension for storing information about files and directories. This is discussed in the Handbook chapter on creating and using CDROMs (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/creating-cds.html), specifically the section on Using Data CDROMs (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/creating-cds.html#MOUNTING-CD).

18. I burned a CD under FreeBSD and now I can not read it under any other operating system. Why?

You most likely burned a raw file to your CD, rather than creating an ISO 9660 filesystem. Take a look at the Handbook chapter on creating CDROMs (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/creating-cds.html), particularly the section on burning raw data CDs (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/creating-cds.html#RAWDATA-CD).

19. How can I create an image of a data CD?

This is discussed in the Handbook section on duplicating data CDs (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/creating-cds.html#IMAGING-CD). For more on working with CDROMs, see the Creating CDs Section (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/creating-cds.html) in the Storage chapter in the Handbook.

20. Why can I not `mount` an audio CD?

If you try to mount an audio CD, you will get an error like `cd9660: /dev/acd0c: Invalid argument`. This is because `mount` only works on filesystems. Audio CDs do not have filesystems; they just have data. You need a program that reads audio CDs, such as the `audio/xmcd` port.

21. How do I mount a multi-session CD?

By default, `mount(8)` will attempt to mount the last data track (session) of a CD. If you would like to load an earlier session, you must use the `-s` command line argument. Please see `mount_cd9660(8)` for specific examples.

22. How do I let ordinary users mount floppies, CDROMs and other removable media?

Ordinary users can be permitted to mount devices. Here is how:

1. As `root` set the `sysctl` variable `vfs.usermount` to 1.

```
# sysctl -w vfs.usermount=1
```

2. As `root` assign the appropriate permissions to the block device associated with the removable media.

For example, to allow users to mount the first floppy drive, use:

```
# chmod 666 /dev/fd0
```

To allow users in the group `operator` to mount the CDROM drive, use:

```
# chgrp operator /dev/acd0c
```

```
# chmod 640 /dev/acd0c
```

3. If you are running FreeBSD 5.X or later, you will need to alter `/etc/devfs.conf` to make these changes permanent across reboots.

As `root`, add the necessary lines to `/etc/devfs.conf`. For example, to allow users to mount the first floppy drive add:

```
# Allow all users to mount the floppy disk.
```

```
own      /dev/fd0      root:operator
```

```
perm     /dev/fd0      0666
```

To allow users in the group `operator` to mount the CD-ROM drive add:

```
# Allow members of the group operator to mount CD-ROMs.
```

```
own      /dev/acd0      root:operator
```

```
perm     /dev/acd0      0660
```

4. Finally, add the line `vfs.usermount=1` to the file `/etc/sysctl.conf` so that it is reset at system boot time.

All users can now mount the floppy `/dev/fd0` onto a directory that they own:

```
% mkdir ~/my-mount-point
```

```
% mount -t msdosfs /dev/fd0 ~/my-mount-point
```

Users in group `operator` can now mount the CDROM `/dev/acd0c` onto a directory that they own:

```
% mkdir ~/my-mount-point
```

```
% mount -t cd9660 /dev/acd0c ~/my-mount-point
```

Unmounting the device is simple:

```
% umount ~/my-mount-point
```

Enabling `vfs.usermount`, however, has negative security implications. A better way to access MS-DOS formatted media is to use the `emulators/mttools` package in the ports collection.

Ὁϕιᾱβὺόϕ: The device name used in the previous examples must be changed according to your configuration.

23. The `du` and `df` commands show different amounts of disk space available. What is going on?

You need to understand what `du` and `df` really do. `du` goes through the directory tree, measures how large each file is, and presents the totals. `df` just asks the filesystem how much space it has left. They seem to be the same thing, but a file without a directory entry will affect `df` but not `du`.

When a program is using a file, and you delete the file, the file is not really removed from the filesystem until the program stops using it. The file is immediately deleted from the directory listing, however. You can see this easily enough with a program such as `more`. Assume you have a file large enough that its presence affects the output of `du` and `df`. (Since disks can be so large today, this might be a *very* large file!) If you delete this file while using `more` on it, `more` does not immediately choke and complain that it cannot view the file. The entry is simply removed from the directory so no other program or user can access it. `du` shows that it is gone — it has walked the directory tree and the file is not listed. `df` shows that it is still there, as the filesystem knows that `more` is still using that space. Once you end the `more` session, `du` and `df` will agree.

Note that Soft Updates can delay the freeing of disk space; you might need to wait up to 30 seconds for the change to be visible!

This situation is common on web servers. Many people set up a FreeBSD web server and forget to rotate the log files. The access log fills up `/var`. The new administrator deletes the file, but the system still complains that the partition is full. Stopping and restarting the web server program would free the file, allowing the system to release the disk space. To prevent this from happening, set up `newsyslog(8)`.

24. How can I add more swap space?

In the Configuration and Tuning

(http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/config-tuning.html) section of the Handbook, you will find a section (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/adding-swap-space.html) describing how to do this.

25. Why does FreeBSD see my disk as smaller than the manufacturer says it is?

Disk manufacturers calculate gigabytes as a billion bytes each, whereas FreeBSD calculates them as 1,073,741,824 bytes each. This explains why, for example, FreeBSD's boot messages will report a disk that supposedly has 80GB as holding 76319MB.

Also note that FreeBSD will (by default) reserve 8% of the disk space.

26. How is it possible for a partition to be more than 100% full?

A portion of each UFS partition (8%, by default) is reserved for use by the operating system and the `root` user. `df(1)` does not count that space when calculating the `Capacity` column, so it can exceed 100%. Also, you will notice that the `Blocks` column is always greater than the sum of the `Used` and `Avail` columns, usually by a factor of 8%.

For more details, look up the `-m` option in `tuneufs(8)`.

Ερώτημα 10 Άλλα: ἄλλοι ὅροι

1. Where are the system start-up configuration files?

The primary configuration file is `/etc/defaults/rc.conf` (see `rc.conf(5)`) System startup scripts such as `/etc/rc` and `/etc/rc.d` (see `rc(8)`) just include this file. *Do not edit this file!* Instead, if there is any entry in `/etc/defaults/rc.conf` that you want to change, you should copy the line into `/etc/rc.conf` and change it there.

For example, if you wish to start `named`, the included DNS server, all you need to do is:

```
# echo named_enable="YES" >> /etc/rc.conf
```

To start up local services, place shell scripts in the `/usr/local/etc/rc.d` directory. These shell scripts should be set executable, and end with a `.sh`.

2. How do I add a user easily?

Use the `adduser(8)` command, or the `pw(8)` command for more complicated situations.

To remove the user, use the `rmuser(8)` command or, if necessary, `pw(8)`.

3. Why do I keep getting messages like `root: not found` after editing my crontab file?

This is normally caused by editing the system crontab (`/etc/crontab`) and then using `crontab(1)` to install it:

```
# crontab /etc/crontab
```

This is not the correct way to do things. The system crontab has a different format to the per-user crontabs which `crontab(1)` updates (the `crontab(5)` manual page explains the differences in more detail).

If this is what you did, the extra crontab is simply a copy of `/etc/crontab` in the wrong format it. Delete it with the command:

```
# crontab -r
```

Next time, when you edit `/etc/crontab`, you should not do anything to inform `cron(8)` of the changes, since it will notice them automatically.

If you want something to be run once per day, week, or month, it is probably better to add shell scripts `/usr/local/etc/periodic`, and let the `periodic(8)` command run from the system cron schedule it with the other periodic system tasks.

The actual reason for the error is that the system crontab has an extra field, specifying which user to run the command as. In the default system crontab provided with FreeBSD, this is `root` for all entries. When this crontab is used as the `root` user's crontab (which is *not* the same as the system crontab), `cron(8)` assumes the string `root` is the first word of the command to execute, but no such command exists.

4. Why do I get the error, you are not in the correct group to su root when I try to su to root?

This is a security feature. In order to su to root (or any other account with superuser privileges), you must be in the wheel group. If this feature were not there, anybody with an account on a system who also found out root's password would be able to gain superuser level access to the system. With this feature, this is not strictly true; su(1) will prevent them from even trying to enter the password if they are not in wheel.

To allow someone to su to root, simply put them in the wheel group.

5. I made a mistake in rc.conf, or another startup file, and now I cannot edit it because the filesystem is read-only. What should I do?

When you get the prompt to enter the shell pathname, simply press ENTER, and run mount / to re-mount the root filesystem in read/write mode. You may also need to run mount -a -t ufs to mount the filesystem where your favorite editor is defined. If your favorite editor is on a network filesystem, you will need to either configure the network manually before you can mount network filesystems, or use an editor which resides on a local filesystem, such as ed(1).

If you intend to use a full screen editor such as vi(1) or emacs(1), you may also need to run export TERM=cons25 so that these editors can load the correct data from the termcap(5) database.

Once you have performed these steps, you can edit /etc/rc.conf as you usually would to fix the syntax error. The error message displayed immediately after the kernel boot messages should tell you the number of the line in the file which is at fault.

6. Why am I having trouble setting up my printer?

Please have a look at the Handbook entry on printing. It should cover most of your problem. See the Handbook entry on printing (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/printing.html).

Some printers require a host-based driver to do any kind of printing. These so-called “WinPrinters” are not natively supported by FreeBSD. If your printer does not work in DOS or Windows NT 4.0, it is probably a WinPrinter. Your only hope of getting one of these to work is to check if the print/pnm2ppa port supports it.

7. How can I correct the keyboard mappings for my system?

Please see the Handbook section on using localization (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/using-localization.html), specifically the section on console setup (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/using-localization.html#SETTING-CONSOLE).

8. Why do I get messages like: unknown: <PNP0303> can't assign resources on boot?

The following is an excerpt from a post to the freebsd-current mailing list.

The “can't assign resources” messages indicate that the devices are legacy ISA devices for which a non-PnP-aware driver is compiled into the kernel. These include devices such as keyboard controllers, the programmable interrupt controller chip, and several other bits of standard infrastructure. The resources cannot be assigned because there is already a driver using those addresses.

—Garrett Wollman <wollman@FreeBSD.org>, 24 April 2001

9. Why can I not get user quotas to work properly?

1. It is possible that your kernel is not configured to use quotas. If this is the case, you will need to add the following line to your kernel configuration file and recompile:

```
options QUOTA
```

Please read the Handbook entry on quotas

(http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/quotas.html) for full details.

2. Do not turn on quotas on /.
 3. Put the quota file on the filesystem that the quotas are to be enforced on, i.e.:

Filesystem	Quota file
/usr	/usr/admin/quotas
/home	/home/admin/quotas
...	...

10. Does FreeBSD support System V IPC primitives?

Yes, FreeBSD supports System V-style IPC, including shared memory, messages and semaphores, in the GENERIC kernel. In a custom kernel, enable this support by adding the following lines to your kernel config.

```
options    SYSVSHM        # enable shared memory
options    SYSVSEM        # enable for semaphores
options    SYSVMSG        # enable for messaging
```

Recompile and install your kernel.

11. What other mail-server software can I use instead of Sendmail?

Sendmail (<http://www.sendmail.org/>) is the default mail-server software for FreeBSD, but you can easily replace it with one of the other MTA (for instance, an MTA installed from the ports).

There are various alternative MTAs in the ports tree already, with mail/exim, mail/postfix, mail/qmail, and mail/zmailer being some of the most popular choices.

Diversity is nice, and the fact that you have many different mail-servers to choose from is considered a good thing; therefore try to avoid asking questions like “Is Sendmail better than Qmail?” in the mailing lists. If you do feel like asking, first check the mailing list archives. The advantages and disadvantages of each and every one of the available MTAs have already been discussed a few times.

12. I have forgotten the root password! What do I do?

Do not panic! Restart the system, type **boot -s** at the Boot: prompt to enter Single User mode. At the question about the shell to use, hit ENTER. You will be dropped to a # prompt. Enter **mount -u /** to remount your root filesystem read/write, then run **mount -a** to remount all the filesystems. Run **passwd root** to change the root password then run **exit(1)** to continue booting.

13. How do I keep **Control+Alt+Delete** from rebooting the system?

If you are using syscons (the default console driver) build and install a new kernel with the line:

```
options SC_DISABLE_REBOOT
```

in the configuration file. If you use the PCVT console driver, use the following kernel configuration line instead.

This can also be done by setting the following sysctl which does not require a reboot or kernel recompile:

```
# sysctl hw.syscons.kbd_reboot=0
```

```
options PCVT_CTRL_ALT_DEL
```

14. How do I reformat DOS text files to UNIX ones?

Use this perl command:

```
% perl -i.bak -npe 's/\r\n/\n/g' file ...
```

file is the file(s) to process. The modification is done in-place, with the original file stored with a .bak extension.

Alternatively you can use the tr(1) command:

```
% tr -d '\r' < dos-text-file > unix-file
```

dos-text-file is the file containing DOS text while *unix-file* will contain the converted output. This can be quite a bit faster than using perl.

15. How do I kill processes by name?

Use killall(1).

16. Why is su bugging me about not being in root's ACL?

The error comes from the Kerberos distributed authentication system. The problem is not fatal but annoying. You can either run su with the -K option, or uninstall Kerberos as described in the next question.

17. How do I uninstall Kerberos?

To remove Kerberos from the system, reinstall the bin distribution for the release you are running. If you have the CDROM, you can mount the cd (we will assume on /cdrom) and run

```
# cd /cdrom/bin
# ./install.sh
```

Alternately, you can remove all MAKE_KERBEROS options from /etc/make.conf and rebuild world.

18. What happened to /dev/MAKEDEV?

FreeBSD 5.X and beyond use the devfs(8) device-on-demand system. Device drivers automatically create new device nodes as they are needed, obsoleting /dev/MAKEDEV.

If you are running FreeBSD 4.X or earlier and /dev/MAKEDEV is missing, then you really do have a problem. Grab a copy from the system source code, probably in /usr/src/etc/MAKEDEV.

19. How do I add pseudoterminals to the system?

If you have lots of telnet, ssh, X, or screen users, you will probably run out of pseudoterminals. Here is how to add more:

1. Build and install a new kernel with the line

```
pseudo-device pty 256
```

in the configuration file.

2. Run the commands

```
# cd /dev
# sh MAKEDEV pty{1,2,3,4,5,6,7}
```

to make 256 device nodes for the new terminals.

3. Edit /etc/ttys and add lines for each of the 256 terminals. They should match the form of the existing entries, i.e. they look like

```
ttymc none network
```

The order of the letter designations is `ttym[pqrsPQRS][0-9a-v]`, using a regular expression.

4. Reboot the system with the new kernel and you are ready to go.

20. Why can I not create the snd0 device?

There is no `snd` device. The name is used as a shorthand for the various devices that make up the FreeBSD sound driver, such as `mixer`, `sequencer`, and `dsp`.

To create these devices you should

```
# cd /dev
# sh MAKEDEV snd0
```

Όχι! Βùόç: You can omit this step if you are running FreeBSD 5.0-RELEASE or newer with devfs(5) enabled.

21. How do I re-read `/etc/rc.conf` and re-start `/etc/rc` without a reboot?

Go into single user mode and then back to multi user mode.

On the console do:

```
# shutdown now
(Note: without -r or -h)

# return
# exit
```

22. I tried to update my system to the latest -STABLE, but got -BETAx, -RC or -PRERELEASE! What is going on?

Short answer: it is just a name. RC stands for “Release Candidate”. It signifies that a release is imminent. In FreeBSD, -PRERELEASE is typically synonymous with the code freeze before a release. (For some releases, the -BETA label was used in the same way as -PRERELEASE.)

Long answer: FreeBSD derives its releases from one of two places. Major, dot-zero, releases, such as 4.0-RELEASE and 5.0-RELEASE, are branched from the head of the development stream, commonly referred to as -CURRENT. Minor releases, such as 4.1-RELEASE or 5.2-RELEASE, have been snapshots of the active -STABLE branch. Starting with 4.3-RELEASE, each release also now has its own branch which can be tracked by people requiring an extremely conservative rate of development (typically only security advisories).

When a release is about to be made, the branch from which it will be derived from has to undergo a certain process. Part of this process is a code freeze. When a code freeze is initiated, the name of the branch is changed to reflect that it is about to become a release. For example, if the branch used to be called 4.5-STABLE, its name will be changed to 4.6-PRERELEASE to signify the code freeze and signify that extra pre-release testing should be happening. Bug fixes can still be committed to be part of the release. When the source code is in shape for the release the name will be changed to 4.6-RC to signify that a release is about to be made from it. Once in the RC stage, only the most critical bugs found can be fixed. Once the release (4.6-RELEASE in this example) and release branch have been made, the branch will be renamed to 4.6-STABLE.

For more information on version numbers and the various CVS branches, refer to the Release Engineering (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/articles/releng/article.html) article.

23. I tried to install a new kernel, and the `chflags` failed. How do I get around this?

Short answer: You are probably at security level greater than 0. Reboot directly to single user mode to install the kernel.

Long answer: FreeBSD disallows changing system flags at security levels greater than 0. You can check your security level with the command:

```
# sysctl kern.securelevel
```

You cannot lower the security level; you have to boot to single mode to install the kernel, or change the security level in `/etc/rc.conf` then reboot. See the `init(8)` manual page for details on `securelevel`, and see `/etc/defaults/rc.conf` and the `rc.conf(5)` manual page for more information on `rc.conf`.

24. I cannot change the time on my system by more than one second! How do I get around this?

Short answer: You are probably at security level greater than 1. Reboot directly to single user mode to change the date.

Long answer: FreeBSD disallows changing the time by more than one second at security levels greater than 1. You can check your security level with the command:

```
# sysctl kern.securelevel
```

You cannot lower the security level; you have to boot to single mode to change the date, or change the security level in `/etc/rc.conf` then reboot. See the `init(8)` manual page for details on `securelevel`, and see `/etc/defaults/rc.conf` and the `rc.conf(5)` manual page for more information on `rc.conf`.

25. Why is `rpc.statd` using 256 megabytes of memory?

No, there is no memory leak, and it is not using 256 Mbytes of memory. For convenience, `rpc.statd` maps an obscene amount of memory into its address space. There is nothing terribly wrong with this from a technical standpoint; it just throws off things like `top(1)` and `ps(1)`.

`rpc.statd(8)` maps its status file (resident on `/var`) into its address space; to save worrying about remapping it later when it needs to grow, it maps it with a generous size. This is very evident from the source code, where one can see that the length argument to `mmap(2)` is `0x10000000`, or one sixteenth of the address space on an IA32, or exactly 256MB.

26. Why can I not unset the `schg` file flag?

You are running at an elevated (i.e., greater than 0) `securelevel`. Lower the `securelevel` and try again. For more information, see the FAQ entry on `securelevel` and the `init(8)` manual page.

27. Why does SSH authentication through `.shosts` not work by default in recent versions of FreeBSD?

The reason why `.shosts` authentication does not work by default in more recent versions of FreeBSD is because `ssh(1)` is not installed `suid root` by default. To “fix” this, you can do one of the following:

- As a permanent fix, set `ENABLE_SUID_SSH` to `true` in `/etc/make.conf` and rebuild `ssh` (or run `make world`).
- As a temporary fix, change the mode on `/usr/bin/ssh` to `4555` by running `chmod 4555 /usr/bin/ssh` as `root`. Then add `ENABLE_SUID_SSH= true` to `/etc/make.conf` so the change takes effect the next time `make world` is run.

28. What is `vnlr`?

`vnlr` flushes and frees `vnodes` when the system hits the `kern.maxvnodes` limit. This kernel thread sits mostly idle, and only activates if you have a huge amount of RAM and are accessing tens of thousands of tiny files.

29. What do the various memory states displayed by `top` mean?

- **Active:** pages recently statistically used.
- **Inactive:** pages recently statistically unused.

- **Cache:** (most often) pages that have percolated from inactive to a status where they maintain their data, but can often be immediately reused (either with their old association, or reused with a new association.) There can be certain immediate transitions from `active` to `cache` state if the page is known to be clean (unmodified), but that transition is a matter of policy, depending upon the algorithm choice of the VM system maintainer.
- **Free:** pages without data content, and can be immediately used in certain circumstances where cache pages might be ineligible. Free pages can be reused at interrupt or process state.
- **Wired:** pages that are fixed into memory, usually for kernel purposes, but also sometimes for special use in processes.

Pages are most often written to disk (sort of a VM sync) when they are in the inactive state, but active pages can also be synced (but requires the availability of certain CPU features.) This depends upon the CPU tracking of the modified bit being available, and in certain situations there can be an advantage for a block of VM pages to be synced, whether they are active or inactive. In most common cases, it is best to think of the inactive queue to be a queue of relatively unused pages that might or might not be in the process of being written to disk. Cached pages are already synced, not mapped, but available for immediate process use with their old association or with a new association. Free pages are available at interrupt level, but cached or free pages can be used at process state for reuse. Cache pages are not adequately locked to be available at interrupt level.

There are some other flags (e.g., busy flag or busy count) that might modify some of the rules that I described.

30. How much free memory is available?

There are a couple of kinds of “free memory”. One kind is the amount of memory immediately available without paging anything else out. That is approximately the size of cache queue + size of free queue (with a derating factor, depending upon system tuning.) Another kind of “free memory” is the total amount of VM space. That can be complex, but is dependent upon the amount of swap space and memory. Other kinds of “free memory” descriptions are also possible, but it is relatively useless to define these, but rather it is important to make sure that the paging rate is kept low, and to avoid running out of swap space.

31. What is `/var/empty`? I can not delete it!

`/var/empty` is a directory that the `sshd(8)` program uses when performing privilege separation. The `/var/empty` directory is empty, owned by `root` and has the `schg` flag set.

Although it is not recommended to delete this directory, to do so you will need to unset the `schg` flag first. See the `chflags(1)` manual page for more information (and bear in mind the answer to the question on unsetting the `schg` flag).

3. Will my existing applications run with the Xorg suite?

The Xorg software is written to the same X11R6 specification that XFree86 is, so basic applications should work unchanged. A few lesser-used protocols have been deprecated (XIE, PEX, and `lbxproxy`), but in the first two cases, the FreeBSD port of XFree86 did not support them either.

4. Why did the X projects split, anyway?

The answer to this question is outside the scope of this FAQ. Note that there are voluminous postings in various mailing list archives on the Internet; please use your favorite search engine to investigate the history instead of asking this question on the FreeBSD mailing lists. It may even be the case that only the participants will ever know for certain.

5. Why did FreeBSD choose to go with the Xorg ports by default?

The Xorg developers claim that their goal is to release more often and incorporate new features more quickly. If they are able to do so, this will be very attractive. Also, their software still uses the traditional X license, while XFree86 is now using their modified one.

Ὁριζήσιμος: This decision is still controversial. Only time will tell which implementation proves technically superior. Each FreeBSD user should decide which they prefer.

6. I want to run X, how do I go about it?

If you would like to add X to an existing installation, you should use the `x11/xorg` meta-port, which will build and install all the necessary components.

Then read and follow the documentation on the `xorgconfig(1)` tool, which assists you in configuring Xorg for your particular graphics card/mouse/etc. You may also wish to examine the `xorgcfg(1)` tool, which provides a graphical interface to the X configuration process.

For further information, read the X11 (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/x11.html) section of the FreeBSD Handbook.

You may also wish to investigate the Xaccel server. See the section on Xi Graphics for more details.

7. I tried to run X, but I get an `KDENABIO failed (Operation not permitted)` error when I type `startx`. What do I do now?

Your system is probably running at a raised `securelevel`. It is not possible to start X at a raised `securelevel` because X requires write access to `/dev/io`. For more information, see at the `init(8)` manual page.

So the question is what else you should do instead, and you basically have two choices: set your `securelevel` back down to zero (usually from `/etc/rc.conf`), or run `xdm(1)` at boot time (before the `securelevel` is raised).

See Α: 14. for more information about running `xdm(1)` at boot time.

8. Why does my mouse not work with X?

If you are using `syscons` (the default console driver), you can configure FreeBSD to support a mouse pointer on each virtual screen. In order to avoid conflicting with X, `syscons` supports a virtual device called `/dev/sysmouse`. All mouse events received from the real mouse device are written to the `sysmouse` device via `moused`. If you wish to use your mouse on one or more virtual consoles, *and* use X, see § 4. and set up `moused`.

Then edit `/etc/X11/xorg.conf` and make sure you have the following lines:

```
Section "InputDevice"
    Option          "Protocol" "SysMouse"
    Option          "Device"   "/dev/sysmouse"
    . . . . .
```

Some people prefer to use `/dev/mouse` under X. To make this work, `/dev/mouse` should be linked to `/dev/sysmouse` (see `sysmouse(4)`):

```
# cd /dev
# rm -f mouse
# ln -s sysmouse mouse
```

9. My mouse has a fancy wheel. Can I use it in X?

Yes.

You need to tell X that you have a 5 button mouse. To do this, simply add the lines `Buttons 5` and `ZAxisMapping 4 5` to the “InputDevice” section of `/etc/X11/xorg.conf`. For example, you might have the following “InputDevice” section in `/etc/X11/xorg.conf`.

Εἰσαγωγή 11-1. “InputDevice” Section for Wheeled Mouse in Xorg configuration file

```
Section "InputDevice"
    Identifier      "Mouse1"
    Driver          "mouse"
    Option          "Protocol" "auto"
    Option          "Device"   "/dev/sysmouse"
    Option          "Buttons"  "5"
    Option          "ZAxisMapping" "4 5"
EndSection
```

Εἰσαγωγή 11-2. “.emacs” example for naive page scrolling with Wheeled Mouse (optional)

```
;; wheel mouse
(global-set-key [mouse-4] 'scroll-down)
(global-set-key [mouse-5] 'scroll-up)
```

10. How do I use remote X displays?

For security reasons, the default setting is to not allow a machine to remotely open a window.

To enable this feature, simply start **X** with the optional `-listen_tcp` argument:

```
% startx -listen_tcp
```

11. Why do X Window menus and dialog boxes not work right?

Try turning off the **Num Lock** key.

If your **Num Lock** key is on by default at boot-time, you may add the following line in the `Keyboard` section of the `/etc/X11/xorg.conf` file.

```
# Let the server do the NumLock processing. This should only be
# required when using pre-R6 clients
    ServerNumLock
```

12. What is a virtual console and how do I make more?

Virtual consoles, put simply, enable you to have several simultaneous sessions on the same machine without doing anything complicated like setting up a network or running X.

When the system starts, it will display a login prompt on the monitor after displaying all the boot messages. You can then type in your login name and password and start working (or playing!) on the first virtual console.

At some point, you will probably wish to start another session, perhaps to look at documentation for a program you are running or to read your mail while waiting for an FTP transfer to finish. Just do **Alt+F2** (hold down the **Alt** key and press the **F2** key), and you will find a login prompt waiting for you on the second “virtual console”! When you want to go back to the original session, do **Alt+F1**.

The default FreeBSD installation has eight virtual consoles enabled. **Alt+F1**, **Alt+F2**, **Alt+F3**, and so on will switch between these virtual consoles.

To enable more of them, edit `/etc/ttys` (see `ttys(5)`) and add entries for `ttyv4` to `ttyvc` after the comment on “Virtual terminals”:

```
# Edit the existing entry for ttyv3 in /etc/ttys and change
# "off" to "on".
ttyv3  "/usr/libexec/getty Pc"          cons25  on secure
ttyv4  "/usr/libexec/getty Pc"          cons25  on secure
ttyv5  "/usr/libexec/getty Pc"          cons25  on secure
ttyv6  "/usr/libexec/getty Pc"          cons25  on secure
ttyv7  "/usr/libexec/getty Pc"          cons25  on secure
ttyv8  "/usr/libexec/getty Pc"          cons25  on secure
ttyv9  "/usr/libexec/getty Pc"          cons25  on secure
ttyva  "/usr/libexec/getty Pc"          cons25  on secure
ttyvb  "/usr/libexec/getty Pc"          cons25  on secure
```

Use as many or as few as you want. The more virtual terminals you have, the more resources that are used; this can be important if you have 8MB RAM or less. You may also want to change the `secure` to `insecure`.

Ὁῖαῖοῖεῖ: If you want to run an X server you *must* leave at least one virtual terminal unused (or turned off) for it to use. That is to say that if you want to have a login prompt pop up for all twelve of your Alt-function keys, you are out of luck - you can only do this for eleven of them if you also want to run an X server on the same machine.

The easiest way to disable a console is by turning it off. For example, if you had the full 12 terminal allocation mentioned above and you wanted to run X, you would change settings for virtual terminal 12 from:

```
ttymb "/usr/libexec/getty Pc"          cons25  on  secure
```

to:

```
ttymb "/usr/libexec/getty Pc"          cons25  off secure
```

If your keyboard has only ten function keys, you would end up with:

```
ttv9  "/usr/libexec/getty Pc"          cons25  off secure
ttva  "/usr/libexec/getty Pc"          cons25  off secure
ttvb  "/usr/libexec/getty Pc"          cons25  off secure
```

(You could also just delete these lines.)

Next, the easiest (and cleanest) way to activate the virtual consoles is to reboot. However, if you really do not want to reboot, you can just shut down the X Window system and execute (as `root`):

```
# kill -HUP 1
```

It is imperative that you completely shut down X Window if it is running, before running this command. If you do not, your system will probably appear to hang/lock up after executing the kill command.

13. How do I access the virtual consoles from X?

Use **Ctrl+Alt+F_n** to switch back to a virtual console. **Ctrl+Alt+F1** would return you to the first virtual console.

Once you are back to a text console, you can then use **Alt+F_n** as normal to move between them.

To return to the X session, you must switch to the virtual console running X. If you invoked X from the command line, (e.g., using `startx`) then the X session will attach to the next unused virtual console, not the text console from which it was invoked. If you have eight active virtual terminals then X will be running on the ninth, and you would use **Alt+F9** to return.

14. How do I start XDM on boot?

There are two schools of thought on how to start xdm(1). One school starts xdm from `/etc/ttys` (see `ttys(5)`) using the supplied example, while the other simply runs xdm from `rc.local` (see `rc(8)`) or from a `X.sh` script in `/usr/local/etc/rc.d`. Both are equally valid, and one may work in situations where the other does not. In both cases the result is the same: X will pop up a graphical login: prompt.

The `ttys` method has the advantage of documenting which vty X will start on and passing the responsibility of restarting the X server on logout to `init`. The `rc.local` method makes it easy to kill xdm if there is a problem starting the X server.

If loaded from `rc.local`, `xm` should be started without any arguments (i.e., as a daemon). `xm` must start AFTER `getty` runs, or else `getty` and `xm` will conflict, locking out the console. The best way around this is to have the script sleep 10 seconds or so then launch `xm`.

If you are to start `xm` from `/etc/ttys`, there still is a chance of conflict between `xm` and `getty(8)`. One way to avoid this is to add the `vt` number in the `/usr/local/lib/X11/xm/Xservers` file.

```
:0 local /usr/local/bin/X vt4
```

The above example will direct the X server to run in `/dev/ttyv3`. Note the number is offset by one. The X server counts the `vt`y from one, whereas the FreeBSD kernel numbers the `vt`y from zero.

15. Why do I get `Couldn't open console` when I run `xconsole`?

If you start X with `startx`, the permissions on `/dev/console` will *not* get changed, resulting in things like `xterm -C` and `xconsole` not working.

This is because of the way console permissions are set by default. On a multi-user system, one does not necessarily want just any user to be able to write on the system console. For users who are logging directly onto a machine with a VTY, the `fbtab(5)` file exists to solve such problems.

In a nutshell, make sure an uncommented line of the form

```
/dev/ttyv0 0600 /dev/console
```

is in `/etc/fbtab` (see `fbtab(5)`) and it will ensure that whomever logs in on `/dev/ttyv0` will own the console.

16. Before, I was able to run XFree86 as a regular user. Why does it now say that I must be `root`?

All X servers need to be run as `root` in order to get direct access to your video hardware. Older versions of XFree86 ($\leq 3.3.6$) installed all bundled servers to be automatically run as `root` (`setuid` to `root`). This is obviously a security hazard because X servers are large, complicated programs. Newer versions of XFree86 do not install the servers `setuid` to `root` for just this reason.

Obviously, running an X server as the `root` user is not acceptable, nor a good idea security-wise. There are two ways to be able to use X as a regular user. The first is to use `xm` or another display manager (e.g., `kdm`); the second is to use the `Xwrapper`.

`xm` is a daemon that handles graphical logins. It is usually started at boot time, and is responsible for authenticating users and starting their sessions; it is essentially the graphical counterpart of `getty(8)` and `login(1)`. For more information on `xm` see the XFree86 documentation (<http://www.xfree86.org/sos/resources.html>), and the the FAQ entry on it.

`Xwrapper` is the X server wrapper; it is a small utility to enable one to manually run an X server while maintaining reasonable safety. It performs some sanity checks on the command line arguments given, and if they pass, runs the appropriate X server. If you do not want to run a display manager for whatever reason, this is for you. If you have installed the complete ports collection, you can find the port in `/usr/ports/x11/wrapper`.

17. Why does my PS/2 mouse misbehave under X?

Your mouse and the mouse driver may have somewhat become out of synchronization.

In rare cases the driver may erroneously report synchronization problem and you may see the kernel message:

```
psmintr: out of sync (xxxx != yyyy)
```

and notice that your mouse does not work properly.

If this happens, disable the synchronization check code by setting the driver flags for the PS/2 mouse driver to 0x100. Enter *UserConfig* by giving the `-c` option at the boot prompt:

```
boot: -c
```

Then, in the *UserConfig* command line, type:

```
UserConfig> flags psm0 0x100
UserConfig> quit
```

18. Why does my PS/2 mouse from MouseSystems not work?

There have been some reports that certain model of PS/2 mouse from MouseSystems works only if it is put into the “high resolution” mode. Otherwise, the mouse cursor may jump to the upper-left corner of the screen every so often.

Specify the flags 0x04 to the PS/2 mouse driver to put the mouse into the high resolution mode. Enter *UserConfig* by giving the `-c` option at the boot prompt:

```
boot: -c
```

Then, in the *UserConfig* command line, type:

```
UserConfig> flags psm0 0x04
UserConfig> quit
```

See the previous section for another possible cause of mouse problems.

19. I want to install different X server.

FreeBSD versions prior 5.3 will use the default **XFree86 4.X**, while latter versions will default to **Xorg**. If you want to run a different X11 implementation than the default one, add the following line to `/etc/make.conf`, (if you do not have this file, create it):

```
X_WINDOW_SYSTEM=          xorg
```

This variable may be set to `xorg`, `xfree86-4`, or `xfree86-3`.

20. How do I reverse the mouse buttons?

Run the command `xmodmap -e "pointer = 3 2 1"` from your `.xinitrc` or `.xsession`.

21. How do I install a splash screen and where do I find them?

FreeBSD have a feature to allow the display of “splash” screens during the boot messages. The splash screens currently must be a 256 color bitmap (`*.BMP`) or ZSoft PCX (`*.PCX`) file. In addition, they must have a resolution of

320x200 or less to work on standard VGA adapters. If you compile VESA support into your kernel, then you can use larger bitmaps up to 1024x768. The actual VESA support can either be compiled directly into the kernel with the VESA kernel config option or by loading the VESA kld module during bootup.

To use a splash screen, you need to modify the startup files that control the boot process for FreeBSD.

You need to create a `/boot/loader.rc` file that contains the following lines:

```
include /boot/loader.4th
start
```

and a `/boot/loader.conf` that contains the following:

```
splash_bmp_load="YES"
bitmap_load="YES"
```

This assumes you are using `/boot/splash.bmp` for your splash screen. If you would rather use a PCX file, copy it to `/boot/splash.pcx`, create a `/boot/loader.rc` as instructed above, and create a `/boot/loader.conf` that contains:

```
splash_pcx_load="YES"
bitmap_load="YES"
bitmap_name="/boot/splash.pcx"
```

Now all you need is a splash screen. For that you can surf on over to the gallery at <http://www.baldwin.cx/splash/>.

22. Can I use the Windows keys on my keyboard in X?

Yes. All you need to do is use `xmodmap(1)` to define what function you wish them to perform.

Assuming all “Windows” keyboards are standard then the keycodes for the 3 keys are

- 115 - Windows key, between the left-hand Ctrl and Alt keys
- 116 - Windows key, to the right of the **AltGr** key
- 117 - **Menu** key, to the left of the right-hand **Ctrl** key

To have the left Windows key print a comma, try this.

```
# xmodmap -e "keycode 115 = comma"
```

You will probably have to re-start your window manager to see the result.

To have the Windows key-mappings enabled automatically every time you start X either put the `xmodmap` commands in your `~/.xinitrc` file or, preferably, create a file `~/.xmodmaprc` and include the `xmodmap` options, one per line, then add the line

```
xmodmap $HOME/.xmodmaprc
```

to your `~/.xinitrc`.

For example, you could map the 3 keys to be **F13**, **F14**, and **F15**, respectively. This would make it easy to map them to useful functions within applications or your window manager, as demonstrated further down.

To do this put the following in `~/.xmodmaprc`.

```
keycode 115 = F13
keycode 116 = F14
keycode 117 = F15
```

If you use `fvwm2`, for example, you could map the keys so that **F13** iconifies (or de-iconifies) the window the cursor is in, **F14** brings the window the cursor is in to the front or, if it is already at the front, pushes it to the back, and **F15** pops up the main Workplace (application) menu even if the cursor is not on the desktop, which is useful if you do not have any part of the desktop visible (and the logo on the key matches its functionality).

The following entries in `~/ .fvwmrc` implement the aforementioned setup:

Key F13	FTIWS	A	Iconify
Key F14	FTIWS	A	RaiseLower
Key F15	A	A	Menu Workplace Nop

23. How can I get 3D hardware acceleration for OpenGL®?

The availability of 3D acceleration depends on the version of XFree86 or Xorg that you are using and the type of video chip you have. If you have an NVIDIA chip, you can use the binary drivers provided for FreeBSD on the Drivers (<http://www.nvidia.com/content/drivers/drivers.asp>) section of their website. For other cards with XFree86-4 or Xorg, including the Matrox G200/G400, ATI Rage 128/Radeon, and 3dfx Voodoo 3, 4, 5, and Banshee, information on hardware acceleration is available on the XFree86-4 Direct Rendering on FreeBSD (<http://people.FreeBSD.org/~anholt/dri/>) page.

ΕὰöÛëáéï 12 Äéêôýùόç

1. Where can I get information on “diskless booting”?

“Diskless booting” means that the FreeBSD box is booted over a network, and reads the necessary files from a server instead of its hard disk. For full details, please read the Handbook entry on diskless booting (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/network-diskless.html)

2. Can a FreeBSD box be used as a dedicated network router?

Yes. Please see the Handbook entry on advanced networking (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/advanced-networking.html), specifically the section on routing and gateways (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/network-routing.html).

3. Can I connect my Windows box to the Internet via FreeBSD?

Typically, people who ask this question have two PCs at home, one with FreeBSD and one with some version of Windows the idea is to use the FreeBSD box to connect to the Internet and then be able to access the Internet from the Windows box through the FreeBSD box. This is really just a special case of the previous question and works perfectly well.

If you are using dialup to connect to the Internet user-mode ppp(8) contains a `-nat` option. If you run ppp(8) with the `-nat` option, set `gateway_enable` to `YES` in `/etc/rc.conf`, and configure your Windows machine correctly, this should work fine. For more information, please see the ppp(8) manual page or the Handbook entry on user PPP (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/userppp.html).

If you are using kernel-mode PPP or have an Ethernet connection to the Internet, you need to use natd(8). Please look at the natd (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/network-natd.html) section of the Handbook for a tutorial.

4. Does FreeBSD support SLIP and PPP?

Yes. See the manual pages for slattach(8), sliplogin(8), ppp(8), and pppd(8). ppp(8) and pppd(8) provide support for both incoming and outgoing connections, while sliplogin(8) deals exclusively with incoming connections, and slattach(8) deals exclusively with outgoing connections.

For more information on how to use these, please see the Handbook chapter on PPP and SLIP (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/ppp-and-slip.html).

If you only have access to the Internet through a “shell account”, you may want to have a look at the `net/slirp` package. It can provide you with (limited) access to services such as ftp and http direct from your local machine.

5. Does FreeBSD support NAT or Masquerading?

Yes. If you want to use NAT over a user PPP connection, please see the Handbook entry on user PPP (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/userppp.html). If you want to use NAT over some other sort of network connection, please look at the natd (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/network-natd.html) section of the Handbook.

6. How do I connect two FreeBSD systems over a parallel line using PLIP?

Please see the PLIP section (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/network-plip.html) of the Handbook.

7. Why can I not create a `/dev/ed0` device?

Because they are not necessary. In the Berkeley networking framework, network interfaces are only directly accessible by kernel code. Please see the `/etc/rc.network` file and the manual pages for the various network programs mentioned there for more information. If this leaves you totally confused, then you should pick up a book describing network administration on another BSD-related operating system; with few significant exceptions, administering networking on FreeBSD is basically the same as on SunOS™ 4.0 or Ultrix.

8. How can I set up Ethernet aliases?

If the alias is on the same subnet as an address already configured on the interface, then add `netmask 0xffffffff` to your `ifconfig(8)` command-line, as in the following:

```
# ifconfig ed0 alias 192.0.2.2 netmask 0xffffffff
```

Otherwise, just specify the network address and netmask as usual:

```
# ifconfig ed0 alias 172.16.141.5 netmask 0xfffff00
```

9. How do I get my 3C503 to use the other network port?

If you want to use the other ports, you will have to specify an additional parameter on the `ifconfig(8)` command line. The default port is `link0`. To use the AUI port instead of the BNC one, use `link2`. These flags should be specified using the `ifconfig_*` variables in `/etc/rc.conf` (see `rc.conf(5)`).

10. Why am I having trouble with NFS and FreeBSD?

Certain PC network cards are better than others (to put it mildly) and can sometimes cause problems with network intensive applications like NFS.

See the Handbook entry on NFS

(http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/network-nfs.html) for more information on this topic.

11. Why can I not NFS-mount from a Linux box?

Some versions of the Linux NFS code only accept mount requests from a privileged port; try

```
# mount -o -P linuxbox:/blah /mnt
```

12. Why can I not NFS-mount from a Sun box?

Sun™ workstations running SunOS 4.X only accept mount requests from a privileged port; try

```
# mount -o -P sunbox:/blah /mnt
```

13. Why does mountd keep telling me it can't change attributes and that I have a bad exports list on my FreeBSD NFS server?

The most frequent problem is not understanding the correct format of `/etc/exports`. Please review `exports(5)` and the NFS (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/network-nfs.html) entry in the Handbook, especially the section on configuring NFS (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/network-nfs.html#CONFIGURING-NFS).

14. Why am I having problems talking PPP to NeXTStep machines?

Try disabling the TCP extensions in `/etc/rc.conf` (see `rc.conf(5)`) by changing the following variable to NO:

```
tcp_extensions=NO
```

Xylogic's Annex boxes are also broken in this regard and you must use the above change to connect through them.

15. How do I enable IP multicast support?

FreeBSD supports multicast host operations by default. If you want your box to run as a multicast router, you need to recompile your kernel with the `MROUTING` option and run `mrouted(8)`. FreeBSD will start `mrouted(8)` at boot time if the flag `mrouted_enable` is set to "YES" in `/etc/rc.conf`.

MBONE tools are available in their own ports category, `mbone` (<http://www.FreeBSD.org/ports/mbone.html>). If you are looking for the conference tools `vic` and `vat`, look there!

16. Which network cards are based on the DEC PCI chipset?

Here is a list compiled by Glen Foster <gfooster@driver.nsta.org>, with some more modern additions:

Ενότητα 12-1. Network cards based on the DEC PCI chipset

Vendor	Model
ASUS	PCI-L101-TB
Accton	ENI1203
Cogent	EM960PCI
Compex	ENET32-PCI
D-Link	DE-530
Dayna	DP1203, DP2100
DEC	DE435, DE450
Danpex	EN-9400P3
JCIS	Condor JC1260
Linksys	EtherPCI
Mylex	LNP101
SMC	EtherPower 10/100 (Model 9332)
SMC	EtherPower (Model 8432)

Vendor	Model
TopWare	TE-3500P
Znyx (2.2.x)	ZX312, ZX314, ZX342, ZX345, ZX346, ZX348
Znyx (3.x)	ZX345Q, ZX346Q, ZX348Q, ZX412Q, ZX414, ZX442, ZX444, ZX474, ZX478, ZX212, ZX214 (10mbps/hd)

17. Why do I have to use the FQDN for hosts on my site?

You will probably find that the host is actually in a different domain; for example, if you are in `foo.example.org` and you wish to reach a host called `mumble` in the `example.org` domain, you will have to refer to it by the fully-qualified domain name, `mumble.example.org`, instead of just `mumble`.

Traditionally, this was allowed by BSD BIND resolvers. However the current version of **bind** (see `named(8)`) that ships with FreeBSD no longer provides default abbreviations for non-fully qualified domain names other than the domain you are in. So an unqualified host `mumble` must either be found as `mumble.foo.example.org`, or it will be searched for in the root domain.

This is different from the previous behavior, where the search continued across `mumble.example.org`, and `mumble.edu`. Have a look at RFC 1535 for why this was considered bad practice, or even a security hole.

As a good workaround, you can place the line

```
search foo.example.org example.org
```

instead of the previous

```
domain foo.example.org
```

into your `/etc/resolv.conf` file (see `resolv.conf(5)`). However, make sure that the search order does not go beyond the “boundary between local and public administration”, as RFC 1535 calls it.

18. Why do I get an error, `Permission denied`, for all networking operations?

If you have compiled your kernel with the `IPFIREWALL` option, you need to be aware that the default policy is to deny all packets that are not explicitly allowed.

If you had unintentionally misconfigured your system for firewalling, you can restore network operability by typing the following while logged in as `root`:

```
# ipfw add 65534 allow all from any to any
```

You can also set `firewall_type="open"` in `/etc/rc.conf`.

For further information on configuring a FreeBSD firewall, see the Handbook chapter (http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/firewalls.html).

19. How much overhead does IPFW incur?

Please see the Handbook’s Firewalls

(http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/firewalls.html) section, specifically the section on

IPFW Overhead & Optimization

(http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/firewalls.html#IPFW-OVERHEAD).

20. Why is my `ipfw` “fwd” rule to redirect a service to another machine not working?

Possibly because you want to do network address translation (NAT) and not just forward packets. A “fwd” rule does exactly what it says; it forwards packets. It does not actually change the data inside the packet. Say we have a rule like:

```
01000 fwd 10.0.0.1 from any to foo 21
```

When a packet with a destination address of `foo` arrives at the machine with this rule, the packet is forwarded to `10.0.0.1`, but it still has the destination address of `foo`! The destination address of the packet is *not* changed to `10.0.0.1`. Most machines would probably drop a packet that they receive with a destination address that is not their own. Therefore, using a “fwd” rule does not often work the way the user expects. This behavior is a feature and not a bug.

See the FAQ about redirecting services, the `natd(8)` manual, or one of the several port redirecting utilities in the ports collection (<http://www.FreeBSD.org/ports/index.html>) for a correct way to do this.

21. How can I redirect service requests from one machine to another?

You can redirect FTP (and other service) request with the `socket` package, available in the ports tree in category “sysutils”. Simply replace the service’s command line to call `socket` instead, like so:

```
ftp stream tcp nowait nobody /usr/local/bin/socket socket ftp.example.com ftp
```

where `ftp.example.com` and `ftp` are the host and port to redirect to, respectively.

22. Where can I get a bandwidth management tool?

There are three bandwidth management tools available for FreeBSD. `dummynet(4)` is integrated into FreeBSD as part of `ipfw(4)`. ALTQ (<http://www.csl.sony.co.jp/person/kjc/programs.html>) is available for free on FreeBSD 4.X and has been integrated into FreeBSD 5.X as part of `pf(4)`. Bandwidth Manager from Emerging Technologies (<http://www.etinc.com/>) is a commercial product.

23. Why do I get `/dev/bpf0: device not configured`?

You are running a program that requires the Berkeley Packet Filter (`bpf(4)`), but it is not in your kernel. Add this to your kernel config file and build a new kernel:

```
pseudo-device bpf          # Berkeley Packet Filter
```

On FreeBSD 4.X and earlier, you must also create the device node. After rebooting, go to the `/dev` directory and run:

```
# sh MAKEDEV bpf0
```

Please see the Handbook entry on device nodes

(http://www.FreeBSD.org/doc/el_GR.ISO8859-7/books/handbook/kernelconfig-nodes.html) for more information on managing devices.

24. How do I mount a disk from a Windows machine that is on my network, like `smbmount` in Linux?

Use the **SMBFS** toolset. It includes a set of kernel modifications and a set of userland programs. The programs and information are available as `net/smbfs` in the ports collection, or in the base system as of 4.5-RELEASE and later.

25. What are these messages about “icmp-response bandwidth limit 300/200 pps” in my log files?

This is the kernel telling you that some activity is provoking it to send more ICMP or TCP reset (RST) responses than it thinks it should. ICMP responses are often generated as a result of attempted connections to unused UDP ports. TCP resets are generated as a result of attempted connections to unopened TCP ports. Among others, these are the kinds of activities which may cause these messages:

- Brute-force denial of service (DoS) attacks (as opposed to single-packet attacks which exploit a specific vulnerability).
- Port scans which attempt to connect to a large number of ports (as opposed to only trying a few well-known ports).

The first number in the message tells you how many packets the kernel would have sent if the limit was not in place, and the second number tells you the limit. You can control the limit using the `net.inet.icmp.icmplim` sysctl variable like this, where 300 is the limit in packets per second:

```
# sysctl -w net.inet.icmp.icmplim=300
```

If you do not want to see messages about this in your log files, but you still want the kernel to do response limiting, you can use the `net.inet.icmp.icmplim_output` sysctl variable to disable the output like this:

```
# sysctl -w net.inet.icmp.icmplim_output=0
```

Finally, if you want to disable response limiting, you can set the `net.inet.icmp.icmplim` sysctl variable (see above for an example) to 0. Disabling response limiting is discouraged for the reasons listed above.

26. What are these `arp: unknown hardware address format error` messages?

This means that some device on your local Ethernet is using a MAC address in a format that FreeBSD does not recognize. This is probably caused by someone experimenting with an Ethernet card somewhere else on the network. You will see this most commonly on cable modem networks. It is harmless, and should not affect the performance of your FreeBSD machine.

27. I have just installed CVSup but trying to execute it produces errors. What is wrong?

First, see if the error message you are receiving is like the one shown below.

```
/usr/libexec/ld-elf.so.1: Shared object "libXaw.so.6" not found
```

Errors like these are caused by installing the `net/cvsup` port on a machine which does not have the **XFree86** suite. If you want to use the GUI included with **CVSup** you will need to install **XFree86** now. Alternatively if you just wish to use **CVSup** from a command line you should delete the package previously installed. Then install the `net/cvsup-without-gui` port. This is covered in more detail in the CVSup section (http://www.FreeBSD.org/doc/en_US.ISO8859-1/books/handbook/cvsup.html) of the Handbook.

Chapter 13 Sandboxing

1. What is a sandbox?

“Sandbox” is a security term. It can mean two things:

- A process which is placed inside a set of virtual walls that are designed to prevent someone who breaks into the process from being able to break into the wider system.

The process is said to be able to “play” inside the walls. That is, nothing the process does in regards to executing code is supposed to be able to breach the walls so you do not have to do a detailed audit of its code to be able to say certain things about its security.

The walls might be a `userid`, for example. This is the definition used in the `security(7)` and `named(8)` man pages.

Take the `ntalk` service, for example (see `/etc/inetd.conf`). This service used to run as `userid root`. Now it runs as `userid tty`. The `tty` user is a sandbox designed to make it more difficult for someone who has successfully hacked into the system via `ntalk` from being able to hack beyond that user id.

- A process which is placed inside a simulation of the machine. This is more hard-core. Basically it means that someone who is able to break into the process may believe that he can break into the wider machine but is, in fact, only breaking into a simulation of that machine and not modifying any real data.

The most common way to accomplish this is to build a simulated environment in a subdirectory and then run the processes in that directory `chroot'd` (i.e. `/` for that process is this directory, not the real `/` of the system).

Another common use is to mount an underlying filesystem read-only and then create a filesystem layer on top of it that gives a process a seemingly writeable view into that filesystem. The process may believe it is able to write to those files, but only the process sees the effects - other processes in the system do not, necessarily.

An attempt is made to make this sort of sandbox so transparent that the user (or hacker) does not realize that he is sitting in it.

UNIX implements two core sandboxes. One is at the process level, and one is at the `userid` level.

Every UNIX process is completely firewalled off from every other UNIX process. One process cannot modify the address space of another. This is unlike Windows where a process can easily overwrite the address space of any other, leading to a crash.

A UNIX process is owned by a particular `userid`. If the `userid` is not the `root` user, it serves to firewall the process off from processes owned by other users. The `userid` is also used to firewall off on-disk data.

2. What is `securelevel`?

The `securelevel` is a security mechanism implemented in the kernel. Basically, when the `securelevel` is positive, the kernel restricts certain tasks; not even the superuser (i.e., `root`) is allowed to do them. At the time of this writing, the `securelevel` mechanism is capable of, among other things, limiting the ability to,

- unset certain file flags, such as `schg` (the system immutable flag),
- write to kernel memory via `/dev/mem` and `/dev/kmem`,
- load kernel modules, and
- alter firewall rules.

To check the status of the `securelevel` on a running system, simply execute the following command:

```
# sysctl kern.securelevel
```

The output will contain the name of the `sysctl(8)` variable (in this case, `kern.securelevel`) and a number. The latter is the current value of the `securelevel`. If it is positive (i.e., greater than 0), at least some of the `securelevel`'s protections are enabled.

You cannot lower the `securelevel` of a running system; being able to do that would defeat its purpose. If you need to do a task that requires that the `securelevel` be non-positive (e.g., an `installworld` or changing the date), you will have to change the `securelevel` setting in `/etc/rc.conf` (you want to look for the `kern_securelevel` and `kern_securelevel_enable` variables) and reboot.

For more information on `securelevel` and the specific things all the levels do, please consult the `init(8)` manual page.

Προειδοποίηση: `Securelevel` is not a silver bullet; it has many known deficiencies. More often than not, it provides a false sense of security.

One of its biggest problems is that in order for it to be at all effective, all files used in the boot process up until the `securelevel` is set must be protected. If an attacker can get the system to execute their code prior to the `securelevel` being set (which happens quite late in the boot process since some things the system must do at start-up cannot be done at an elevated `securelevel`), its protections are invalidated. While this task of protecting all files used in the boot process is not technically impossible, if it is achieved, system maintenance will become a nightmare since one would have to take the system down, at least to single-user mode, to modify a configuration file.

This point and others are often discussed on the mailing lists, particularly the `FreeBSD Security` (http://lists.FreeBSD.org/mailman/listinfo/freebsd-security). Please search the archives here (http://www.FreeBSD.org/search/index.html) for an extensive discussion. Some people are hopeful that `securelevel` will soon go away in favor of a more fine-grained mechanism, but things are still hazy in this respect.

Consider yourself warned.

3. `BIND (named)` is listening on port 53 and some other high-numbered port. What is going on?

`BIND` uses a random high-numbered port for outgoing queries. If you want to use port 53 for outgoing queries, either to get past a firewall or to make yourself feel better, you can try the following in `/etc/namedb/named.conf`:

```
options {
    query-source address * port 53;
};
```

You can replace the `*` with a single IP address if you want to tighten things further.

Congratulations, by the way. It is good practice to read your `sockstat(1)` output and notice odd things!

4. `Sendmail` is listening on port 587 as well as the standard port 25! What is going on?

Recent versions of `Sendmail` support a mail submission feature that runs over port 587. This is not yet widely supported, but is growing in popularity.

5. What is this UID 0 `toor` account? Have I been compromised?

Do not worry. `toor` is an “alternative” superuser account (`toor` is `root` spelt backwards). Previously it was created when the `bash(1)` shell was installed but now it is created by default. It is intended to be used with a non-standard shell so you do not have to change `root`’s default shell. This is important as shells which are not part of the base distribution (for example a shell installed from ports or packages) are likely to be installed in `/usr/local/bin` which, by default, resides on a different filesystem. If `root`’s shell is located in `/usr/local/bin` and `/usr` (or whatever filesystem contains `/usr/local/bin`) is not mounted for some reason, `root` will not be able to log in to fix a problem (although if you reboot into single user mode you will be prompted for the path to a shell).

Some people use `toor` for day-to-day `root` tasks with a non-standard shell, leaving `root`, with a standard shell, for single user mode or emergencies. By default you cannot log in using `toor` as it does not have a password, so log in as `root` and set a password for `toor` if you want to use it.

6. Why is `suidperl` not working properly?

For security reasons, `suidperl` is installed without the `suid` bit by default. The system administrator can enable `suid` behavior with the following command.

```
# chmod u+s /usr/bin/suidperl
```

If you want `suidperl` to be built `suid` during upgrades from source, edit `/etc/make.conf` and add `ENABLE_SUIDPERL=true` before you run `make buildworld`.

